arXiv:1608.07347v4 [math.PR] 13 Feb 2017

LOWER BOUNDS FOR THE SMALLEST SINGULAR VALUE OF STRUCTURED RANDOM MATRICES

NICHOLAS COOK

Abstract. We obtain lower tail estimates for the smallest singular value of random matrices with independent but non-identically distributed entries. Specifically, we consider $n \times n$ matrices with complex entries of the form

$$M = A \circ X + B = (a_{ij}\xi_{ij} + b_{ij})$$

where $X = (\xi_{ij})$ has iid centered entries of unit variance and A and B are fixed matrices. In our main result we obtain polynomial bounds on the smallest singular value of M for the case that A has bounded (possibly zero) entries, and $B = Z\sqrt{n}$ where Z is a diagonal matrix with entries bounded away from zero. As a byproduct of our methods we can also handle general perturbations B under additional hypotheses on A, which translate to connectivity hypotheses on an associated graph. In particular, we extend a result of Rudelson and Zeitouni for Gaussian matrices to allow for general entry distributions satisfying some moment hypotheses. Our proofs make use of tools which (to our knowledge) were previously unexploited in random matrix theory, in particular Szemerédi's Regularity Lemma, and a version of the Restricted Invertibility Theorem due to Spielman and Srivastava.

Contents

1. Introduction 2
   1.1. Background 2
   1.2. A general class of non-iid matrices 5
   1.3. Previous results 7
   1.4. New results 9
   1.5. Ideas of the proof 13
   1.6. Organization of the paper 16
   1.7. Notation 16
2. Preliminaries 17
   2.1. Partitioning and discretizing the sphere 17
   2.2. Anti-concentration for scalar random walks 18
   2.3. Anti-concentration for the image of a fixed vector 20

Date: February 15, 2017.
2010 *Mathematics Subject Classification.* Primary: 60B20, secondary: 15B52.
The author was partially supported by NSF grant DMS-1266164.

3. Invertibility from connectivity: Compressible vectors 22
3.1. Highly compressible vectors 24
3.2. An entropy reduction lemma 25
3.3. Broadly connected profile: Proof of Proposition 3.1 28
3.4. General profile: Proof of Proposition 3.2 31
4. Invertibility from connectivity: Incompressible vectors 35
4.1. Proof of Theorem 1.13 36
4.2. Proof of Theorem 1.25 37
5. Invertibility under diagonal perturbation: Proof of main theorem 40
5.1. Preliminary Tools 40
5.2. Control on the operator norm 42
5.3. Decomposition of the standard deviation profile 45
5.4. High level proof of Theorem 1.19 48
5.5. Proof of Lemma 5.10 52
5.6. Proof of Lemma 5.11 54
Appendix A. Proofs of anti-concentration lemmas 55
A.1. Proof of Lemma 2.5 56
A.2. Proof of Lemma 2.7 57
References 58

1. Introduction

1.1. **Background.** Recall that the singular values of an $n \times n$ matrix M with complex entries are the eigenvalues of $\sqrt{M^*M}$, which we arrange in non-increasing order:

$$s_1(M) \geq \cdots \geq s_n(M) \geq 0.$$

We have $s_1(M) = \|M\|$ (throughout we write $\|\cdot\|$ for the $\ell_2^n \to \ell_2^n$ operator norm), and M is invertible if and only if $s_n(M) > 0$, in which case we have $s_n(M) = \|M^{-1}\|^{-1}$. We (informally) say that M is "well-invertible" if $s_n(M)$ is well-separated from zero.

The largest and smallest singular values of random matrices with independent entries have been intensely studied, in part due to applications in theoretical computer science. Motivated by their work on the first electronic computers, von Neumann and Goldstine sought upper bounds on the condition number $\kappa(M) = s_1(M)/s_n(M)$ of a large matrix M with iid entries [vNG47] (at the time, "large" was on the order $n \approx 100$); see Chapter 2 of Edelman's PhD thesis [Ede89] for further background and history. More recently, bounds on the condition number of non-centered random matrices have been important in the theory of *smoothed analysis of algorithms* developed by Spielman and Teng [SST06]. The smallest singular value of shifted random matrices has also received attention due to its

connection with proving convergence of the empirical spectral distribution; see [TV08] and the survey [BC12].

Much is known about the largest singular value for random matrices with independent entries. First we review the iid case: we denote by $X = X_n$ an $n \times n$ matrix whose entries ξ_{ij} are iid copies of a centered complex random variable with unit variance, and refer to such X as an "iid matrix". From the works [BSY88,YBK88] it is known that $\frac{1}{\sqrt{n}} s_1(X_n) \in (2-\varepsilon, 2+\varepsilon)$ with probability tending to one as $n \to \infty$ for any fixed $\varepsilon > 0$ (if we further assume a coupling wherein each matrix X_n is the top-left $n \times n$ corner of an infinite array of iid variables then this can be strengthened to the almost sure convergence $\frac{1}{\sqrt{n}} s_1(X_n) \to 2$).

In connection with problems in computer science and the theory of Banach spaces there has been considerable interest in obtaining non-asymptotic bounds. Seginer [Seg00] proved that if X is an $n \times m$ matrix with iid centered entries, then

$$\mathbb{E}\, s_1(X) \ll \mathbb{E} \max_{i \in [n]} \sqrt{\sum_{j \in [m]} \xi_{ij}^2} + \mathbb{E} \max_{j \in [m]} \sqrt{\sum_{i \in [n]} \xi_{ij}^2} \tag{1.1}$$

(see Section 1.7 for our conventions on asymptotic notation). Latała [Lat05] subsequently gave a bound (Theorem 5.7 below) that holds for matrices with independent but not necessarily identically distributed entries, with entries having finite fourth moment. The noncommutative Kintchine inequality [LP86, LPP91] and the method of "matrix concentration" [Oli10, Tro12] also provide non-asymptotic bounds for the spectral norm of centered random matrices, though these fail to be sharp for the case of the iid matrix X. Recent works of Bandeira–van Handel [BvH16] and van Handel [vH] provide bounds which are sharp in many cases. In [vH] van Handel records the following conjecture (which he attributes to Latała). The conjecture remains open as of this writing.

Conjecture 1.1. *The bound* (1.1) *holds for arbitrary centered Gaussian matrices.*

As the two terms on the right hand side in (1.1) are each lower bounds for the left hand side, the conjecture states that the two sides are equivalent up to a universal constant factor. One might hope that Conjecture 1.1 extends to matrices with more general entry distributions, but Seginer [Seg00] gives an example of a block diagonal matrix with ± 1 Bernoulli entries for which (1.1) fails. See [BvH16, Remark 4.8] for further discussion.

The picture is far less complete for the smallest singular value of random matrices; however, recent years have seen much progress for the case of the iid matrix X. The limiting distribution of $\sqrt{n} s_n(X)$ was obtained by Edelman for the case of Gaussian entries [Ede88], and this law was shown by Tao and Vu to hold for all iid matrices with entries ξ_{ij} having a sufficiently large finite moment [TV10a].

Quantitative lower tail estimates for $s_n(X)$ proved to be considerably more challenging than bounding the operator norm. The first breakthrough was made by Rudelson [Rud08], who showed that if X has iid real-valued sub-Gaussian

entries, that is

$$\forall p \geq 1, \qquad (\mathbb{E}\,|\xi|^p)^{1/p} \leq C\sqrt{p} \tag{1.2}$$

for some constant $C < \infty$, then

$$\mathbb{P}\Big(s_n(X) \leq tn^{-3/2}\Big) \ll t + n^{-1/2} \tag{1.3}$$

for all $t \geq 0$, where the implied constant depends on the constant C.

Around the same time, in [TV09] Tao and Vu obtained bounds of the form

$$\mathbb{P}\Big(s_n(X) \leq n^{-\beta}\Big) \ll n^{-\alpha} \tag{1.4}$$

for any fixed $\alpha > 0$ and β sufficiently large depending on α, for the case that the entries of X take values in $\{-1, 0, 1\}$. They extended (1.4) to uncentered matrices with general entry distributions having finite second moment in [TV08] (see Theorem 1.7 below), which was instrumental for their proof of a strong version of the celebrated circular law for the limiting spectral distribution of $\frac{1}{\sqrt{n}}X$.

The results in [TV09] are based on powerful results on the Littlewood–Offord problem, which is to establish anti-concentration (or "small ball") bounds for random walks of the form

$$W(v) = \sum_{j=1}^{n} \xi_j v_j \tag{1.5}$$

with iid variables ξ_j and a fixed vector of weights $v \in \mathbb{C}^n$. Such random walks arise in the proof of (1.4) as the dot product between a row of X and a fixed vector v. Starting with work of Halász in the 1970s, a close connection was observed between anti-concentration properties of $W(v)$ and the presence of additive structure in the components of v (see also independent work of Arak from the 1980s [Ara80, Ara81, GEZ16]). Using tools from additive combinatorics, Tao and Vu established a strong quantitative version of this so-called *Littlewood–Offord phenomenon* in the form of an inverse theorem, showing that if $W(v)$ concentrates in some small region with high probability (of size $n^{-O(1)}$) then the components of v have some special additive structure (most of them are contained in a *generalized arithmetic progression* of controlled complexity). (In the earlier work [TV07] they had used Freiman's theorem from additive combinatorics to prove a different inverse-type result, which they used to bound the singularity probability for random Bernoulli matrices.)

Motivated by these developments, in [RV08b] Rudelson and Vershynin found a different way to quantify the additive structure of a vector v – called the *essential least common denominator* – which they used to obtain the following improvement of (1.3) for matrices with sub-Gaussian entries:

$$\mathbb{P}\Big(s_n(X) \leq tn^{-1/2}\Big) \ll t + e^{-cn}. \tag{1.6}$$

This estimate is optimal up to the implied constant and $c > 0$ (both of which depend on the constant C in (1.2)). Thus, they achieved a stronger bound than (1.4) under stronger assumptions.

See the survey [NV13] for additional background and results in the Littlewood–Offord theory. We highlight a few subsequent developments on lower tail bounds for the smallest singular value in Section 1.3 below. We also mention that there has also been work on *upper tail bounds* for the smallest singular value – see in particular [RV08a, NV16] – but we do not consider this problem further in the present work.

1.2. **A general class of non-iid matrices.** In this paper we are concerned with bounds for the smallest singular value of random matrices with independent but non-identically distributed entries. The following definition allows us to quantify the dependence of our bounds on the distribution of the matrix entries.

Definition 1.2 (Spread random variable)**.** Let ξ be a complex random variable and let $\kappa \geq 1$. We say that ξ is κ*-spread* if

$$\operatorname{Var}\left[\, \xi \, \mathbb{1}(|\xi - \mathbb{E}\,\xi| \leq \kappa) \,\right] \geq \frac{1}{\kappa}. \tag{1.7}$$

Remark 1.3. It follows from the monotone convergence theorem that any random variable ξ with non-zero second moment is κ-spread for some $\kappa < \infty$. Furthermore, if ξ is centered with unit variance and finite pth moment μ_p for some $p > 2$, then it is routine to verify that ξ is κ-spread with $\kappa = 3(3\mu_p^p)^{1/(p-2)}$, say.

Our results concern the following general class of matrices:

Definition 1.4 (Structured random matrix)**.** Let $A = (a_{ij})$ and $B = (b_{ij})$ be deterministic $n \times m$ matrices with $a_{ij} \in [0,1]$ and $b_{ij} \in \mathbb{C}$ for all i, j. Let $X = (\xi_{ij})$ be an $n \times m$ matrix with independent entries, all identically distributed to a complex random variable ξ with mean zero and variance one. Put

$$M = A \circ X + B = (a_{ij}\xi_{ij} + b_{ij})_{i,j=1}^{n} \tag{1.8}$$

where $\circ$ denotes the matrix Hadamard product. We refer to A, B and ξ as the *standard deviation profile*, *mean profile* and *atom variable*, respectively. We denote the L^p norm of the atom variable by

$$\mu_p := (\mathbb{E}\,|\xi|^p)^{1/p}. \tag{1.9}$$

Without loss of generality, we assume throughout that ξ is κ_0-spread for some fixed $\kappa_0 \geq 1$.

(While all of our results are for square matrices, we give the definition for the general rectangular case as we will often need to consider rectangular submatrices in the proofs.)

Remark 1.5. The assumption that the entries of M are shifted scalings of random variables ξ_{ij} having a common distribution is made for convenience, as it allows us to access some standard anti-concentration estimates (see Section 2.2). We expect the proofs can be modified to cover general matrices with independent entries having specified means and variances (possibly with additional moment hypotheses), but we do not pursue this here.

As a concrete example one can consider a centered non-Hermitian band matrix, where one sets $a_{ij} \equiv 0$ for $|i-j|$ exceeding some bandwidth parameter $w \in [n-1]$ – see Corollary 1.17.

The singular value distributions for structured random matrices have been studied in connection with wireless MIMO networks [TV04,HLN07]. The limiting spectral distributions and spectral radius for certain structured random matrices have been used to model the dynamical properties of neural networks [RA06, ARS15]. In a forthcoming paper with Hachem, Najim and Renfrew [CHNR], the limiting spectral distribution is determined for a general class of centered structured random matrices. That work required bounds on the smallest singular value for shifts of centered matrices by scalar multiples of the identity, which was the original motivation for the results in this paper (in particular, Corollary 1.22 below is a key input for the proofs in [CHNR]).

The picture for the smallest singular value of structured random matrices is far less complete than for the largest singular value – in particular, it is not clear to this author what is the correct analogue of Conjecture 1.1. Here we content ourselves with identifying sufficient conditions on the matrices A, B and the distribution of ξ for a structured random matrix M to be well-invertible with high probability. Specifically, we seek to address the following:

Question A. Let M be an $n \times n$ random matrix as in Definition 1.4. Under what assumptions on the standard deviation and mean profiles A, B and the distribution of the atom variable ξ do we have

$$\mathbb{P}\left(s_n(M) \leq n^{-\beta}\right) = O(n^{-\alpha}) \tag{1.10}$$

for some constants $\alpha, \beta > 0$?

The case that $B = -z\sqrt{n}I$ for some fixed $z \in \mathbb{C}$ (where I denotes the $n \times n$ identity matrix) is of particular interest for applications to the limiting spectral distribution of centered random matrices. As we shall see in the next subsection, existing results in the literature give lower tail bounds for $s_n(M)$ that are uniform in the shift B under the size constraint $\|B\| = n^{O(1)}$, i.e.

$$\sup_{B \in \mathcal{M}_n(\mathbb{C}):\, \|B\| \leq n^C} \mathbb{P}\left(s_n(A \circ X + B) \leq n^{-\beta}\right) = O(n^{-\alpha}). \tag{1.11}$$

for some constant $C > 0$ (results stated for centered matrices generally extend in a routine manner to allow a perturbation of size $\|B\| = O(\sqrt{n})$). Such bounds can be viewed as matrix analogues of classical anti-concentration (or "small ball") bounds of the form

$$\sup_{z \in \mathbb{C}} \mathbb{P}(|S_n - z| \leq r) = o_{r \to 0}(1) + o_{n \to \infty}(1) \tag{1.12}$$

where S_n is a sequence of scalar random variables (such as the normalized partial sums of an infinite sequence of iid variables). In fact, bounds of the form (1.12) are a central ingredient in the proofs of estimates (1.11). Roughly speaking, the translation invariance of (1.12) causes the uniformity in the shift B in (1.11)

to come for free once one can handle the centered case $B = 0$ (the assumption $\|B\| = n^{O(1)}$ is needed to have some continuity of the map $u \mapsto \|Mu\|$ on the unit sphere in order to apply a discretization argument). In light of this we may pose the following:

Question B. Let M be an $n \times n$ random matrix as in Definition 1.4, and let $\gamma > 0$. Under what assumptions on the standard deviation profile A and the distribution of the atom variable ξ do we have

$$\sup_{B \in \mathcal{M}_n(\mathbb{C}):\, \|B\| \le n^\gamma} \mathbb{P}\left(s_n(M) \le n^{-\beta}\right) = O(n^{-\alpha}) \tag{1.13}$$

for some constants $\alpha, \beta > 0$?

The following simple observation puts a clear limitation on the standard deviation profiles A for which we can expect to have (1.13).

Observation 1.6. *Suppose that $A = (a_{ij})$ has a $k \times m$ submatrix of zeros for some k, m with $k + m > n$. Then $A \circ X$ is singular with probability 1. Thus,* (1.13) *fails (by taking $B = 0$) for any fixed $\alpha, \beta > 0$.*

Theorem 1.13 below (see also Theorem 1.11 for the Gaussian case) shows that the above is in some sense the only obstruction to obtaining (1.13).

1.3. **Previous results.** Before stating our main results on Questions A and B we give an overview of what is currently in the literature.

For the case of a constant standard deviation profile A and essentially arbitrary mean profile B we have the following result of Tao and Vu, which was a key technical component of their proof of the circular law for iid matrices [TV08, TV10b].

Theorem 1.7 (Shifted iid matrix [TV08])**.** *Let X be an $n \times n$ matrix with iid entries $\xi_{ij} \in \mathbb{C}$ having mean zero and variance one. For any $\alpha, \gamma > 0$ there exists $\beta > 0$ such that for any fixed (deterministic) $n \times n$ matrix B with $\|B\| \le n^\gamma$,*

$$\mathbb{P}\left(s_n(X + B) \le n^{-\beta}\right) = O_{\alpha,\gamma}(n^{-\alpha}). \tag{1.14}$$

A stronger version of the above bound was established earlier by Sankar, Spielman and Teng for the case that X has iid standard Gaussian entries [SST06]. For the case that $B = 0$, the bound (1.6) of Rudelson and Vershynin gives the optimal dependence $\beta = \alpha + 1/2$ for the exponents, but requires the stronger assumption that the entries are real-valued and sub-Gaussian (we remark that their proof extends in a routine manner to allow an arbitrary shift B with $\|B\| = O(\sqrt{n})$). Recently, the sub-Gaussian assumption for (1.6) was relaxed by Rebrova and Tikhomirov to only assume a finite second moment [RT].

When the entries of M have bounded density the problem is much simpler. The following is easily obtained by the argument in [BC12, Section 4.4].

Proposition 1.8 (Matrix with entries having bounded density [BC12])**.** *Let M be an $n \times n$ random matrix with independent entries having density on $\mathbb{C}$ or $\mathbb{R}$ uniformly bounded by $\varphi > 1$. For every $\alpha > 0$ there is a $\beta = \beta(\alpha, \varphi) > 0$ such that*

$$\mathbb{P}\left(s_n(M) \leq n^{-\beta}\right) = O(n^{-\alpha}). \tag{1.15}$$

Note that above we make no assumptions on the moments of the entries of M – in particular, they may have heavy tails. The following result of Bordenave and Chafaï (Lemma A.1 in [BC12]) relaxes the hypothesis of continuous distributions from Proposition 1.8 while still allowing for heavy tails, but comes at the cost of a worse probability bound.

Proposition 1.9 (Heavy-tailed matrix with non-degenerate entries [BC12])**.** *Let Y be an $n \times n$ random matrix with independent entries $\eta_{ij} \in \mathbb{C}$. Suppose that for some $p, r, \sigma_0 > 0$ we have that for all $i, j \in [n]$,*

$$\mathbb{P}(|\eta_{ij}| \leq r) \geq p, \qquad \mathrm{Var}(\eta_{ij}\, \mathbb{1}(|\eta_{ij}| \leq r)) \geq \sigma_0^2. \tag{1.16}$$

For any $s \geq 1$, $t \geq 0$, and any fixed $n \times n$ matrix B we have

$$\mathbb{P}\bigg(s_n(Y + B) \leq \frac{t}{\sqrt{n}},\ \|Y + B\| \leq s\bigg) \ll_{p,r,\sigma_0} \sqrt{\log s}\bigg(ts^2 + \frac{1}{\sqrt{n}}\bigg). \tag{1.17}$$

The non-degeneracy conditions (1.16) do not allow for some entries to be deterministic. Litvak and Rivasplata [LR12] obtained a lower tail estimate of the form (1.10) for centered random matrices having a sufficiently small constant proportion of entries equal to zero deterministically. Below we give new results (Theorems 1.13 and 1.25) allowing all but an arbitrarily small (fixed) proportion of entries to be deterministic.

Finally, we recall a theorem of Rudelson and Zeitouni [RZ16] which shows that, for the case of Gaussian matrices, Observation 1.6 is essentially the only obstruction to obtaining (1.13). To state their result we need to set up some graph theoretic notation, which will be used repeatedly throughout the paper.

To a non-negative $n \times m$ matrix $A = (a_{ij})$ we associate a bipartite graph $\Gamma_A = ([n], [m], E_A)$, with $(i, j) \in E_A$ if and only if $a_{ij} > 0$. For a row index $i \in [n]$ we denote by

$$\mathcal{N}_A(i) = \{j \in [m] : a_{ij} > 0\} \tag{1.18}$$

its neighborhood in Γ_A. Thus, the neighborhood of a column index $j \in [m]$ is denoted $\mathcal{N}_{A^\intercal}(j)$. Given sets of row and column indices $I \subset [n], J \subset [m]$, we define the associated *edge count*

$$e_A(I, J) := |\{(i, j) \in [n] \times [m] : a_{ij} > 0\}|. \tag{1.19}$$

We will generally work with the graph that only puts an edge (i, j) when a_{ij} exceeds some fixed cutoff parameter $\sigma_0 > 0$. Thus, we denote by

$$A(\sigma_0) = (a_{ij} 1_{a_{ij} \geq \sigma_0}) \tag{1.20}$$

the matrix which thresholds out entries smaller than σ_0.

Rudelson and Zeitouni obtain control on the smallest singular value of Gaussian matrices whose associated matrix of standard deviations $A = (a_{ij})$ satisfies the following expansion-type condition.

Definition 1.10 (Broad connectivity)**.** Let $A = (a_{ij})$ be an $n \times m$ matrix with non-negative entries. For $I \subset [n]$ and $\delta \in (0,1)$, define the set of δ-*broadly connected* neighbors of I as

$$\mathcal{N}_A^{(\delta)}(I) = \{j \in [m] : |\mathcal{N}_{A^\intercal}(j) \cap I| \geq \delta |I|\}. \tag{1.21}$$

For $\delta, \nu \in (0,1)$, we say that A is (δ,ν)-*broadly connected* if

(1) $|\mathcal{N}_A(i)| \geq \delta m$ for all $i \in [n]$;

(2) $|\mathcal{N}_{A^\intercal}(j)| \geq \delta n$ for all $j \in [m]$;

(3) $|\mathcal{N}_{A^\intercal}^{(\delta)}(J)| \geq \min(n, (1+\nu)|J|)$ for all $J \subset [m]$.

Theorem 1.11 (Gaussian matrix with broadly connected profile [RZ16])**.** *Let G be an $n \times n$ matrix with iid standard real Gaussian entries, and let A be an $n \times n$ matrix with entries $a_{ij} \in [0,1]$ for all i, j. With notation as in (1.20), assume that $A(\sigma_0)$ is (δ,ν)-broadly connected for some $\sigma_0, \delta, \nu \in (0,1)$. Let $K \geq 1$, and let B be a fixed $n \times n$ matrix with $\|B\| \leq K\sqrt{n}$. Then for any $t \geq 0$,*

$$\mathbb{P}\left(s_n(A \circ G + B) \leq tn^{-1/2}\right) \ll K^{O(1)}t + e^{-cn} \tag{1.22}$$

where the implied constants and $c > 0$ depend only on δ, ν, and σ_0.

Note that the assumption of broad connectivity gives us an "epsilon of separation" from the bad example of Observation 1.6. Thus, Theorem 1.11 provides a near-optimal answer to Question B for Gaussian matrices.

Remark 1.12. Since the dependence of the bound (1.22) on the parameters δ and ν is not quantified, Theorem 1.11 only addresses Question B for *dense* standard deviation profiles, i.e. when A has a non-vanishing proportion of large entries. While it would not be difficult to quantify the steps in [RZ16], the resulting dependence on parameters is not likely to be optimal.

1.4. **New results.** Our first result removes the Gaussian assumption from Theorem 1.11, though at the cost of a worse probability bound. Recall the parameter κ_0 from Definition 1.4.

Theorem 1.13 (General matrix with broadly connected profile)**.** *Let $M = A \circ X + B$ be an $n \times n$ matrix as in Definition 1.4, and assume that $A(\sigma_0)$ is (δ,ν)-broadly connected for some $\sigma_0, \delta, \nu \in (0,1)$. Let $K \geq 1$. For any $t \geq 0$,*

$$\mathbb{P}\left(s_n(M) \leq \frac{t}{\sqrt{n}},\ \|M\| \leq K\sqrt{n}\right) \ll t + \frac{1}{\sqrt{n}} \tag{1.23}$$

where the implied constant depends only on K, δ, ν, σ_0 and κ_0.

Remark 1.14. While we have stated no moment assumptions on the atom variable ξ over the standing assumption of unit variance, the restriction to the event $\{\|M\| \le K\sqrt{n}\}$ requires us to assume at least four finite moments to deduce $\mathbb{P}(s_n(M) \le t/\sqrt{n}) \ll t + o(1)$. Here we give a lower tail estimate at the optimal scale $s_n(M) \sim n^{-1/2}$; however, the arguments in this paper can be used to establish a polynomial lower bound on $s_n(M)$ of non-optimal order under larger perturbations B (similar to (1.32) below).

Remark 1.15 (Improving the probability bound). We expect that the probability bound in (1.23) can be improved by making use of the more advanced tools of Littlewood–Offord theory introduced in [TV08, RV08b], though it appears that these tools cannot be applied in a straightforward manner. In the interest of keeping the paper of reasonable length we do not pursue this matter here.

Remark 1.16 (Bounds on moderately small singular values). The methods used to prove Theorem 1.13 together with an idea of Tao and Vu from [TV10b] can be used to give lower bounds of optimal order on "moderately small" singular values, i.e. $s_{n-k}(M)$ with $n^{\varepsilon} \le k \le cn$ for any $\varepsilon > 0$ and a sufficiently small constant $c = c(\kappa_0, \sigma_0, \delta, \nu, K) > 0$. See Theorem 4.5.1 in the author's PhD thesis [Coo16]. Such bounds are of interest for proving convergence of the empirical spectral distribution – see [TV10b, BC12].

In light of Observation 1.6, Theorem 1.13 gives an essentially optimal answer to Question B for *dense* random matrices (see Remark 1.12). It would be interesting to establish a version of this result that allows for only a proportion $o(1)$ of the entries to be random. Indeed, we expect a version of the above theorem to hold when A has density as small $(\log^{O(1)} n)/n$. (Quantifying the dependence on δ, ν in (1.23) would only allow a slight polynomial decay in the density.)

We note that they broad connectivity hypothesis includes many standard deviation profiles of interest, such as band matrices:

Corollary 1.17 (Shifted non-Hermitian band matrices)**.** *Let $M = A \circ X + B$ be an $n \times n$ matrix as in Definition 1.4, and assume that for some fixed $\sigma_0, \varepsilon \in (0,1)$, $a_{ij} \ge \sigma_0$ for all i, j with $\min(|i-j|, n-|i-j|) \le \varepsilon n$. Let $K \ge 1$. Then (1.23) holds for any $t \ge 0$, where the implied constant depends only on K, σ_0, ε and κ_0.*

Remark 1.18. For applications to dynamical systems (as in [RA06, ARS15]) the matrix A encodes a network in which nodes live on a cycle graph of size n and only interact with nodes within distance εn in the graph metric. One can easily modify the argument below to give the same conclusion if we only assume $a_{ij} \ge \sigma_0$ when $|i-j| \le \varepsilon n$, i.e. replacing the cycle with a line graph.

Proof. By conditioning on the entries ξ_{ij} with $\min(|i-j|, n-|i-j|) > \varepsilon n$ and absorbing the corresponding entries of $A \circ X$ into B we may assume the entries of $A(\sigma_0)$ are zero outside the band. By Theorem 1.13 it suffices to show that $A(\sigma_0)$ is $(c\varepsilon, c\varepsilon)$-broadly connected for a sufficiently small constant $c > 0$. For all

$i \in [n]$ we have $|\mathcal{N}_{A(\sigma_0)}(i)|, |\mathcal{N}_{A^{\mathsf{T}}(\sigma_0)}(i)| \geq 2\varepsilon n + 1$, so it only remains to verify the third property in Definition 1.10. For the remainder of the proof we abbreviate $\mathcal{N}(i) = \mathcal{N}_{A(\sigma_0)}(i)$ and $\mathcal{N}^{(\delta)}(J) = \mathcal{N}^{(\delta)}_{A^{\mathsf{T}}(\sigma_0)}(J)$.

Define $f : [n] \to [0,1]$ by $f(i) = |J \cap \mathcal{N}(i)|/(2\varepsilon n + 1)$. By double counting we have

$$\sum_{i=1}^{n} f(i) = |J|. \tag{1.24}$$

On the other hand, we have the discrete derivative bound

$$|f(i) - f(i-1)| \leq 1/(2\varepsilon n + 1) \tag{1.25}$$

(here is where we use the band structure of the matrix). Suppose that

$$|\mathcal{N}^{(\delta)}(J)| = |\{i : f(i) \geq \delta |J|/n\}| < (1+\nu)|J|.$$

From (1.25) we can bound

$$\sum_{i \in \mathcal{N}^{(\delta)}(J)} f(i) \leq (2\varepsilon n + 1)\Big(\frac{\delta |J|}{n} + \frac{1}{2}\Big) + \Big((1+\nu)|J| - (2\varepsilon n + 1)\Big).$$

Together with (1.24) we find

$$\begin{aligned} |J| &= \sum_{i \notin \mathcal{N}^{(\delta)}(J)} f(i) + \sum_{i \in \mathcal{N}^{(\delta)}(J)} f(i) \\ &< |J|(1 + \delta + \nu) - \Big(\frac{1}{2} - \delta\Big) 2\varepsilon n \end{aligned}$$

and taking $\delta = \nu = \varepsilon/4$ leads to a contradiction. □

Having addressed Question B, we now ask whether we can further relax the assumptions on the standard deviation profile A by assuming more about the mean profile B. In particular, can we make assumptions on B that give (1.10) while allowing $A \circ X$ to be singular deterministically (as in Observation 1.6)?

Of course, a trivial example is to take $A = 0$ and B any invertible matrix. Another easy example is to take take B to be *very well-invertible*, with $s_n(B) \geq K\sqrt{n}$ for a large constant $K > 0$ (for instance, take $B = K\sqrt{n}I$, where I is the identity matrix). Indeed, standard estimates for the operator norm of random matrices with centered entries (cf. Section 5.2) give $\|A \circ X\| = O(\sqrt{n})$ with high probability provided the atom variable ξ satisfies some additional moment hypotheses. From the triangle inequality

$$\begin{aligned} s_n(M) &= \inf_{u \in S^{n-1}} \|(A \circ X + B)u\| \\ &\geq \inf_{u \in S^{n-1}} \|Bu\| - \sup_{u \in S^{n-1}} \|(A \circ X)u\| \\ &= s_n(B) - \|A \circ X\| \end{aligned}$$

so $s_n(M) \gg \sqrt{n}$ with high probability if K is sufficiently large.

The problem becomes non-trivial when we allow B to have singular values of size $\varepsilon\sqrt{n}$ for small $\varepsilon > 0$ and A as in Observation 1.6. In this case any proof of a lower tail estimate of the form (1.10) must depart significantly from the proofs of the results in the previous section by making use of arguments which are not translation invariant.

Our main result shows that when the mean profile B is a *diagonal* matrix with smallest entry at least an arbitrarily small (fixed) multiple of $\sqrt{n}$, then we do not need to assume anything further about the standard deviation profile A.

Theorem 1.19 (Main result)**.** *Fix arbitrary $r_0 \in (0, \frac{1}{2}]$, $K_0 \geq 1$, and let Z be a (deterministic) diagonal matrix with diagonal entries $z_1, \dots, z_n \in \mathbb{C}$ satisfying*

$$|z_i| \in [r_0, K_0] \qquad \forall i \in [n]. \tag{1.26}$$

Let M be an $n \times n$ random matrix as in Definition 1.4 with $B = Z\sqrt{n}$, and assume $\mu_{4+\eta} < \infty$ for some fixed $\eta > 0$. There are $\alpha(\eta) > 0$ and $\beta(r_0, \eta, \mu_{4+\eta}) > 0$ such that

$$\mathbb{P}\left(s_n(M) \leq n^{-\beta}\right) = O(n^{-\alpha}) \tag{1.27}$$

where the implied constant depends only on r_0, K_0 and $\mu_{4+\eta}$.

Remark 1.20 (Moment assumption). The assumption of $4 + \eta$ moments is due to our use of a result of Vershynin, Theorem 5.8 below, on the operator norm of products of random matrices. Apart from this, at many points in our argument we use that an $m \times m$ submatrix of M has operator norm $O(\sqrt{m})$ with high probability (assuming m grows with n), which requires at least four finite moments. However, under certain additional assumptions on the standard deviation profile we only need to assume two finite moments – see Remark 5.12.

Remark 1.21 (Dependence of α, β on parameters). The dependence of α on η is of the form $\alpha = c\min(1, \eta)$ for a constant $c > 0$ (which we take to be $1/9$ in the proof). If we were to assume ξ has finite pth moment for a sufficiently large constant p then the proof would allow us to take any fixed $\alpha < 1/2$ in (1.27). The dependence of β on $\nu_{4+\eta}$ and r_0 given by our proof is very bad, of the form

$$\beta = \mathrm{twr}\big(O_\eta(1)\exp((\mu_{4+\eta}/r_0)^{O(1)})\big) \tag{1.28}$$

where $\mathrm{twr}(x)$ is a tower exponential $2^{2^{\cdot^{\cdot^{\cdot^{2}}}}}$ of height x. (The factor $O_\eta(1)$ comes from Vershynin's bound mentioned in the previous remark – we do not know the precise dependence on η, but we expect it is relatively mild.) This is due to our use of Szemerédi's regularity lemma (specifically, a version for directed graphs due to Alon and Shapira – see Lemma 5.2). It would be interesting to obtain a version of Theorem 1.19 with a better dependence of β on the parameters.

As we remarked above, the case of a diagonal mean profile is of special interest for the problem of proving convergence of the empirical spectral distribution of centered random matrices with a variance profile.

Corollary 1.22 (Scalar shift of a centered random matrix)**.** *Let $X = (\xi_{ij})$ be an $n \times n$ matrix whose entries are iid copies of a centered complex random variable ξ having unit variance and $(4+\eta)$-th moment $\mu_{4+\eta} < \infty$ for some fixed $\eta > 0$. Let $A = (a_{ij})$ be a fixed $n \times n$ non-negative matrix with entries uniformly bounded by $\sigma_{\max} < \infty$. Put $Y = \frac{1}{\sqrt{n}} A \circ X$, and fix an arbitrary $z \in \mathbb{C} \setminus \{0\}$. There are constants $\alpha = \alpha(\eta) > 0$ and $\beta = \beta(|z|, \eta, \mu_{4+\eta}, \sigma_{\max}) > 0$ such that*

$$\mathbb{P}\left(s_n(Y - zI) \leq n^{-\beta}\right) = O(n^{-\alpha}) \tag{1.29}$$

where the implied constant depends only on $|z|, \sigma_{\max}$ and $\mu_{4+\eta}$.

While our main motivation was to handle diagonal perturbations of centered random matrices, we conjecture that Theorem 1.19 extends to matrices as in Definition 1.4 with more general mean profiles B:

Conjecture 1.23. *Theorem 1.19 continues to hold for $B \in \mathcal{M}_n(\mathbb{C})$ not necessarily diagonal, where the constraint* (1.26) *is replaced with $\frac{1}{\sqrt{n}} s_i(B) \in [r_0, K_0]$ for all $1 \leq i \leq n$.*

1.5. **Ideas of the proof.** Here we give an informal discussion of the main ideas in the proof of Theorem 1.19.

Regular partitions of graphs. As with Theorem 1.13, the key is to associate the standard deviation profile A with a graph. Since we want the diagonal of M to be preserved under relabeling of vertices will will associate A with a directed graph (digraph) which puts an edge $i \to j$ whenever a_{ij} exceeds some small threshold $\sigma_0 > 0$. Since A has no special connectivity structure *a priori*, we will apply a version of Szemerédi's regularity lemma for digraphs (Lemma 5.2) to partition the vertex set $[n]$ into a bounded number of parts of equal size $I_1, \ldots, I_m$, together with a small set of "bad" vertices I_{bad}, such that for most $(k, l) \in [m]^2$ the subgraph on $I_k \cup I_l$ enjoys certain "pseudorandomness" properties. These properties will not be quite strong enough to control the smallest singular value of the corresponding submatrix M_{I_k, I_l} of M, but we can apply a "cleaning" procedure (as it is called in the extremal combinatorics literature) to remove a small number of bad vertices from each part in the partition (which we add to I_{bad}), after which we will be able to control $s_{\min}(M_{I_k, I_l})$ for most $(k, l) \in [m]^2$. We defer the precise formulation of the pseudorandomness properties and corresponding bound on the smallest singular value to Definition 1.24 and Theorem 1.25 below.

Schur complement formula. The task will then be to lift this control on the invertibility of submatrices to the whole matrix M. The key tool here is the *Schur complement formula* (see Lemma 5.4) which allows us to control the smallest singular value of a block matrix

$$\begin{pmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{pmatrix} \tag{1.30}$$

assuming some control on the smallest singular values of (perturbations of) the diagonal block submatrices M_{11}, M_{22} and on the operator norm of the off-diagonal submatrices M_{12}, M_{21}. The control on the smallest singular value of the whole matrix is somewhat degraded, but this is acceptable as we will only apply Lemma 5.4 a bounded number of times. If we can find a generalized diagonal of "good" block submatrices that are well-invertible under additive perturbations, then after permuting the blocks to lie on the main diagonal we can apply the Schur complement bound along a nested sequence of submatrices partitioned as in (1.30), where M_{11} is a "good" matrix and M_{22} is well-invertible by the induction hypothesis. We remark that the strategy of leveraging properties of a small submatrix using the Schur complement formula was recently applied in a somewhat different manner in [BEYY] to prove the universality of spectral statistics of random Hermitian band matrices.

Decomposition of the reduced digraph. At this point it is best to think of the regular partition $I_1, \dots, I_m$ as inducing a "macroscopic scale" digraph $\mathcal{R} = ([m], E)$ (often called the *reduced digraph* in extremal combinatorics) that puts an edge $(k, l) \in E$ whenever the corresponding submatrix A_{I_k, I_l} is pseudorandom and sufficiently dense. If we can cover the vertices of $\mathcal{R}$ with vertex-disjoint directed cycles, then we will have found a generalized diagonal of submatrices of M with the desired properties, and we can finish with a bounded number of applications of the Schur complement formula as described above.

Of course, it may be the case that $\mathcal{R}$ cannot be covered by disjoint cycles. For instance, if A were to have all ones in the first $n/2$ columns and all zeros in the last $n/2$ columns then roughly half of the vertices of $\mathcal{R}$ would have no incoming edges. This is where we make crucial use of the diagonal perturbation $Z\sqrt{n}$ (indeed, without this perturbation M would be singular in this example). The top left $n/2 \times n/2$ submatrix of M is dense, and we can apply Theorem 1.25 to control its smallest singular vale. The bottom right $n/2 \times n/2$ submatrix is a diagonal matrix with diagonal entries of size at least $r_0\sqrt{n}$, and hence its smallest singular value is at least $r_0\sqrt{n}$. This argument even allows for the bottom right submatrix of A to be nonzero but sufficiently sparse: we can use the triangle inequality and standard bounds on the operator norm of sparse random matrices to argue that the smallest singular value of the bottom right submatrix is still of order $\gg r_0\sqrt{n}$.

We handle the general case as follows. We greedily cover as many of the vertices of $\mathcal{R}$ as we can with disjoint cycles – call this set of vertices $U_{\text{cyc}} \subset [m]$. At this point we have either covered the whole graph (and we are done) or the graph on the remaining vertices U_{free} is cycle-free. This means that the vertices of $\mathcal{R}$ can be relabeled so that its adjacency matrix is upper-triangular on $U_{\text{free}} \times U_{\text{free}}$. Write $J_{\text{cyc}} = \bigcup_{k \in U_{\text{cyc}}} I_k$, $J_{\text{free}} = \bigcup_{k \in U_{\text{free}}} I_k$ and denote the corresponding submatrices of A on the diagonal by $A_{\text{cyc}}, A_{\text{free}}$, and likewise for M. We thus have a relabeling of $[n]$ under which A_{free} is close to upper triangular (there may be some entries of A_{free} below the diagonal of size less than σ_0, or which are contained in a small number of exceptional pairs from the regular partition). Crucially, this relabeling

has preserved the diagonal, so the submatrix M_{free} is a diagonal perturbation of an (almost) upper-triangular random matrix. We then show that such a matrix has smallest singular value of order $\gg_{r_0} \sqrt{n}$ with high probability. With another application of the Schur complement bound we can combine the control on the submatrices $M_{\text{cyc}}, M_{\text{free}}$ (along with standard bounds on the operator norm for the off-diagonal blocks) to conclude the proof. (Actually, the bad set I_{bad} of rows and columns requires some additional arguments, but we do not discuss these here.)

This concludes the high level description of the proof of Theorem 1.19. We only remark that the above partitioning and cleaning procedures will generate various error terms and residual submatrices (such as the vertices in I_{bad}, or the small proportion of pairs (I_k, I_l) which are not sufficiently pseudorandom). As the smallest singular value is notoriously sensitive to perturbations, it will take some care to control these terms. We will use some high-powered tools such as bounds on the operator norm of sparse random matrices and products of random matrices due to Latała and Vershynin – see Section 5.2.

Invertibility from connectivity assumptions. Now we state the specific pseudorandomness condition on a standard deviation profile under which we have good control on the smallest singular value. While "pseudorandom" generally means that the edge distribution in a graph is close to uniform on a range of scales, we will only need control from below on the edge densities (morally speaking, we want the matrix A to be as far as possible from the zero matrix, the most poorly invertible matrix). The following one-sided pseudorandomness condition is taken from the combinatorics literature (see [KS96, Definition 1.6]). The reader should recall the notation introduced in (1.18)–(1.20).

Definition 1.24 (Super-regularity)**.** Let A be an $n \times m$ matrix with non-negative entries. For $\delta, \varepsilon \in (0,1)$, we say that A is (δ, ε)-*super-regular* if the following hold:
(1) $|\mathcal{N}_A(i)| \geq \delta m$ for all $i \in [n]$;
(2) $|\mathcal{N}_{A^\intercal}(j)| \geq \delta n$ for all $j \in [m]$;
(3) $e_A(I, J) \geq \delta |I||J|$ for all $I \subset [n], J \subset [m]$ such that $|I| \geq \varepsilon n$ and $|J| \geq \varepsilon m$.

The reader should compare this condition with Definition 1.10. Conditions (1) and (2) are are the same in both definitions, while it is not hard to see that condition (3) above implies

$$|\mathcal{N}_{A^\intercal}^{(\delta)}(J)| \geq (1-\varepsilon)n \tag{1.31}$$

whenever $|J| \geq \varepsilon n$ (with notation as in (1.21)), which is stronger than condition (3) in Definition 1.10 for such J. On the other hand, conditions (1) and (2) imply that $|\mathcal{N}_{A^\intercal}^{(\sqrt{\delta}/2)}(J)| \geq \frac{1}{2}\delta n$ for any $J \subset [m]$ (see Lemma 3.4), so super-regularity is stronger than broad connectivity for ε, η sufficiently small depending on δ.

Theorem 1.25 (Matrix with super-regular profile)**.** *Let $M = A \circ X + B$ be an $n \times n$ matrix as in Definition 1.4. Assume that $A(\sigma_0)$ (as defined in* (1.20)*)*

is (δ,ε)-super-regular for some $\delta,\sigma_0 \in (0,1)$ and $0 < \varepsilon < c_1\delta\sigma_0^2$ with $c_1 > 0$ a sufficiently small constant. For any $\gamma \geq 1/2$ there exists $\beta = O(\gamma^2)$ such that

$$\mathbb{P}\left(s_n(M) \leq n^{-\beta},\, \|M\| \leq n^{\gamma}\right) \ll \sqrt{\frac{\log n}{n}} \tag{1.32}$$

where the implied constant depends only on γ, δ, σ_0 and κ_0.

Note that Theorem 1.25 allows for a mean profile B of arbitrary polynomial size in operator norm, whereas in Theorem 1.13 we only allowed $\|B\| = O(\sqrt{n})$. The ability to handle such large perturbations will be crucial in the proof of Theorem 1.19, as the iterative application of the Schur complement bound discussed above will lead to perturbations of increasingly large polynomial order.

We defer discussion of the key technical ideas for Theorem 1.13 and Theorem 1.25 to Sections 3 and 4. We only mention here that our proof of Theorem 1.25 makes crucial use of a new "entropy reduction" argument, which allows us to control the event that $\|Mu\|$ is small for some u in certain portions of the sphere S^{n-1} by the event that this holds for some u in a random net of relatively low cardinality. The argument uses an improvement by Spielman and Srivastava [SS12] of the classic Restricted Invertibility Theorem due to Bourgain and Tzafriri [BT87] – see Section 3 for details.

1.6. **Organization of the paper.** The rest of the paper is organized as follows. Sections 2, 3 and 4 are devoted to the proofs of Theorems 1.13 and 1.25. We prove these theorems in parallel as they involve many similar ideas. In Section 2 we collect some standard lemmas on anti-concentration for random walks and products of random matrices with fixed vectors, along with some facts about nets in Euclidean space. In Section 3 we show that random matrices as in Theorems 1.13 and 1.25 are well-invertible over sets of "compressible" vectors in the unit sphere, and in Section 4 we establish control over the complementary set of "incompressible" vectors. Theorem 1.19 is proved in Section 5.

1.7. **Notation.** $\mathcal{M}_{n,m}(\mathbb{C})$ denotes the set of $n \times m$ matrices with complex entries. When $m = n$ we will write $\mathcal{M}_n(\mathbb{C})$. For a matrix $A = (a_{ij}) \in \mathcal{M}_{n,m}(\mathbb{C})$ we will sometimes use the notation $A(i,j) = a_{ij}$. For $I \subset [n], J \subset [m]$, $A_{I,J}$ denotes the $|I| \times |J|$ submatrix with entries indexed by $I \times J$. We abbreviate $A_J := A_{J,J}$.

$\|\cdot\|$ denotes the Euclidean norm when applied to vectors, and the $\ell_2^m \to \ell_2^n$ operator norm when applied to an $n \times m$ matrix. $\|A\|_{\mathrm{HS}}$ denotes the Hilbert–Schmidt (or Frobenius) norm of a matrix A. We will sometimes denote the smallest singular value of a square matrix M by $s_{\min}(M)$ (in situations where M is a submatrix of a larger matrix this will often be clearer than writing the dimension).

We denote the unit sphere in $\mathbb{C}^n$ by S^{n-1}. For $J \subset [n]$, we denote by $\mathbb{C}^J \subset \mathbb{C}^n$ (resp. $S^J \subset S^{n-1}$) the set of vectors (resp. unit vectors) in $\mathbb{C}^n$ supported on J. Given a vector $v \in \mathbb{C}^n$, we denote by $v_J \in \mathbb{C}^n$ the projection of v to the coordinate subspace $\mathbb{C}^J$. For $m \in \mathbb{N}$, $x \in \mathbb{R}$, $\binom{[m]}{x}$ denotes the family of subsets of $[m]$ of size $\lfloor x \rfloor$.

C, c, c', c_0, etc. denote unspecified constants, and their value may change from line to line. They are understood to be absolute if no dependence on parameters is indicated. $f = O(g)$, $f \ll g$, $g \gg f$ are all synonymous and mean that $|f| \leq Cg$ for some absolute constant $C < \infty$. We will occasionally use the notation $f = o(g)$ to mean that $f/g \to 0$ as $n \to \infty$, where the parameter n will be the size of the matrix under consideration (this will only be for the sake of brevity, as all of our arguments are quantitative). We indicate dependence of the implied constant C on parameters with subscripts, e.g. $f \ll_\alpha g$. $f = o_\alpha(g)$ means that $f/g \to 0$ when α is fixed and $n \to \infty$.

When considering a random matrix M as in Definition 1.4, we write

$$\mathcal{F}_{I,J} := \langle \{\xi_{ij}\}_{i \in I, j \in J} \rangle \tag{1.33}$$

for the sigma algebra of events generated by the entries $\{\xi_{ij}\}_{i \in I, j \in J}$ of X. For $I \subset [n]$ we write $\mathbb{P}_I(\cdot)$ for probability conditional on $\mathcal{F}_{[n] \setminus I, [n]}$.

Acknowledgements. The author thanks David Renfrew and Terence Tao for useful conversations, and also thanks David Renfrew for providing helpful comments on a preliminary version of the manuscript.

2. Preliminaries

2.1. **Partitioning and discretizing the sphere.** For the proofs of Theorems 1.13 and 1.25 we make heavy use of ideas and notation developed in [LPRTJ05, LPR$^+$05, Rud08, RV08b] and related ideas from geometric functional analysis. In particular, in order to lower bound

$$s_n(M) = \inf_{u \in S^{n-1}} \|Mu\|$$

we partition the sphere into sets of vectors of different levels of "compressibility", which we presently define, and separately obtain control on the infimum of $\|Mu\|$ over each set.

Recall from Section 1.7 our notation $\mathbb{C}^J \subset \mathbb{C}^m$ for the set of vectors supported on $J \subset [m]$. For a set $T \subset \mathbb{C}^n$ and $\rho > 0$ we write T_ρ for the set of points within Euclidean distance ρ of T. We recall also the following definitions from [RZ16]. For $\theta, \rho \in (0,1)$, we define the set of *compressible vectors*

$$\mathrm{Comp}(\theta, \rho) := S^{m-1} \cap \bigcup_{J \in \binom{[m]}{\theta m}} (\mathbb{C}^J)_\rho \tag{2.1}$$

and the complementary set of *incompressible vectors*

$$\mathrm{Incomp}(\theta, \rho) := S^{m-1} \setminus \mathrm{Comp}(\theta, \rho). \tag{2.2}$$

That is, $\mathrm{Comp}(\theta, \rho)$ is the set of unit vectors within (Euclidean) distance ρ of a vector supported on at most θm coordinates. On the other hand, incompressible vectors enjoy the following property which will lead to good anti-concentration properties for an associated random walk.

Lemma 2.1 (Incompressible vectors are spread, cf. [RV08b, Lemma 3.4])**.** *Fix $\theta, \rho \in (0,1)$ and let $v \in \mathrm{Incomp}(\theta, \rho)$. There is a set $L^+ \subset [m]$ with $|L^+| \geq \theta m$ such that $|v_j| \geq \rho/\sqrt{m}$ for all $j \in L^+$. Moreover, for all $\lambda \geq 1$ there is a set $L \subset [m]$ with $|L| \geq (1 - \frac{1}{\lambda^2})\theta m$ such that for all $j \in L$,*

$$\frac{\rho}{\sqrt{m}} \leq |v_j| \leq \frac{\lambda}{\sqrt{\theta m}}.$$

Proof. Take $L^+ = \{j : |v_j| \geq \rho/\sqrt{m}\}$ and denote $L^- = \{j : |v_j| \leq \lambda/\sqrt{\theta m}\}$. Since v lies a distance at least ρ from any vector supported on at most θm coordinates we must have $|L^+| \geq \theta m$, which gives the first claim. On the other hand, since $v \in S^{m-1}$, by Markov's inequality we have $|L^-| \geq \theta m/\lambda^2$, so taking $L = L^+ \cap L^-$ we have $|L| \geq (1 - \frac{1}{\lambda^2})\theta m$. □

For fixed choices of θ, ρ we informally refer to the coordinates of $v \in \mathrm{Incomp}(\theta, \rho)$ where $|v_j| \geq \rho/\sqrt{n}$ as the *essential support of* v.

Now we recall a standard fact about nets of the sphere of controlled cardinality. For $\rho > 0$, recall that a ρ-net of a set $T \subset \mathbb{C}^m$ is a finite subset $\Sigma \subset T$ such that for all $v \in T$ there exists $v' \in \Sigma$ with $\|v - v'\| \leq \rho$.

Lemma 2.2 (Metric entropy of the sphere)**.** *Let $V \subset \mathbb{C}^m$ be a subspace of (complex) dimension k, let $T \subset V \cap S^{m-1}$, and let $\rho \in (0,1)$. Then T has a ρ-net $\Sigma \subset T$ of cardinality $|\Sigma| \leq (3/\rho)^{2k}$.*

Proof. Let $\Sigma \subset T$ be a ρ-separated (in Euclidean distance) subset that is maximal under set inclusion. It follows from maximality that Σ is a ρ-net of T. Let $\Sigma_{\rho/2}$ denote the $\rho/2$ neighborhood of Σ in V. Noting that $\Sigma_{\rho/2}$ is a disjoint union of k-dimensional Euclidean balls of radius $\rho/2$, we have

$$|\Sigma| c_k (\rho/2)^{2k} \leq \mathrm{vol}_k(\Sigma_{\rho/2}) \leq c_k (1 + \rho/2)^{2k}$$

where vol_k denotes the k-dimensional Lebesgue measure on V and c_k is the volume of the Euclidean unit ball in $\mathbb{C}^k$. The desired bound follows by rearranging. □

2.2. **Anti-concentration for scalar random walks.** In this subsection we collect some standard anti-concentration estimates for scalar random walks, which are perhaps the most central tool for proving that random matrices are (well-)invertible with high probability.

Definition 2.3 (Concentration probability)**.** Let ξ be a complex-valued random variable. For $v \in \mathbb{C}^n$ we let

$$S_\xi(v) = \sum_{j=1}^n \xi_j v_j \tag{2.3}$$

where $\xi_1, \ldots, \xi_n$ are iid copies of ξ. For $r \geq 0$ we define the *concentration probability*

$$p_{\xi,v}(r) = \sup_{z \in \mathbb{C}} \mathbb{P}\big(|S_\xi(v) - z| \leq r\big). \tag{2.4}$$

Throughout this section we operate under the following distributional assumption on ξ.

Definition 2.4 (Controlled second moment, cf. [TV08, Definition 2.2])**.** Let $\kappa \geq 1$. A complex random variable ξ is said to have *κ-controlled second moment* if one has the upper bound

$$\mathbb{E}\,|\xi|^2 \leq \kappa \tag{2.5}$$

(in particular, $|\,\mathbb{E}\,\xi| \leq \kappa^{1/2}$), and the lower bound

$$\mathbb{E}[\mathrm{Re}(z\xi - w)]^2\,\mathbb{1}(|\xi| \leq \kappa) \geq \frac{1}{\kappa}[\mathrm{Re}(z)]^2 \tag{2.6}$$

for all $z \in \mathbb{C}, a \in \mathbb{R}$.

Roughly speaking, a complex random variable ξ has controlled second moment if its distribution has a one-(real-)dimensional marginal with fairly large variance on some compact set. The following is a quantitative version of [TV08, Lemma 2.4], and shows that by multiplying the matrices X and B in Definition 1.4 by a scalar phase (amounting to multiplying M by a phase, which does not affect its singular values) we can assume the atom variable ξ has $O(\kappa_0)$-controlled second moment in all of our proofs with no loss of generality. The proof is deferred to Appendix A.1.

Lemma 2.5. *Let ξ be a centered complex random variable with unit variance, and assume ξ is κ_0-spread for some $\kappa_0 \geq 1$ (see Definition 1.2). Then there exists $\theta \in \mathbb{R}$ such that $e^{i\theta}\xi$ has κ-controlled second moment for some $\kappa = O(\kappa_0)$.*

Below we give two standard bounds on the concentration function $p_{\xi,v}(r)$ when ξ is a κ-controlled random variable and $v \in S^{n-1}$. The first gives a crude constant order bound that is uniform in $v \in S^{n-1}$:

Lemma 2.6 (Crude anti-concentration, cf. [TV10c, Corollary 6.3])**.** *Let ξ be a complex random variable with κ-controlled second moment. There exists $r_0 > 0$ depending only on κ such that $p_{\xi,v}(r_0) \leq 1 - r_0$ for all $v \in S^{n-1}$.*

Note that Lemma 2.6 is sharp for the case that v is a standard basis vector. The following gives an improved bound when v has small ℓ_∞ norm.

Lemma 2.7 (Improved anti-concentration)**.** *Let ξ be a complex random variable that is κ-controlled for some $\kappa > 0$, and let $v \in S^{n-1}$. For all $r \geq 0$,*

$$p_{\xi,v}(r) \ll_\kappa r + \|v\|_\infty. \tag{2.7}$$

Lemma 2.7 can be deduced from the Berry–Esséen theorem (which is the approach taken in [LPRTJ05], for instance), but this would require ξ to have finite third moment, which we do not assume. (Generally speaking, higher moment assumptions should only be necessary to prove concentration bounds as opposed to anti-concentration.) Since we could not locate a proof in the literature for the case that ξ and and the coefficients of v take values in $\mathbb{C}$, we provide a proof in Appendix A.2.

2.3. **Anti-concentration for the image of a fixed vector.** In this subsection we boost the anti-concentration bounds for scalar random variables from the previous sections to anti-concentration for the image of a fixed vector under a random matrix. The following lemma of Rudelson and Vershynin is convenient for this task.

Lemma 2.8 (Tensorization, cf. [RV08b, Lemma 2.2])**.** *Let $\zeta_1, \dots, \zeta_n$ be independent non-negative random variables.*

(a) Suppose that for some $\varepsilon_0, p_0 > 0$ and all $j \in [n]$, $\mathbb{P}(\zeta_j \le \varepsilon_0) \le p_0$. There are $c_1, p_1 \in (0,1)$ depending only on p_0 such that

$$\mathbb{P}\left(\sum_{j=1}^n \zeta_j^2 \le c_1 \varepsilon_0^2 n \right) \le p_1^n. \tag{2.8}$$

(b) Suppose that for some $K, \varepsilon_0 \ge 0$ and all $j \in [n]$, $\mathbb{P}(\zeta_j \le \varepsilon) \le K\varepsilon$ for all $\varepsilon \ge \varepsilon_0$. Then for all $\varepsilon \ge \varepsilon_0$,

$$\mathbb{P}\left(\sum_{j=1}^n \zeta_j^2 \le \varepsilon^2 n \right) \le (CK\varepsilon)^n. \tag{2.9}$$

Proof. We only prove part (a), as part (b) is given in [RV08b] (note that in part (a) we have given more specific dependencies on the parameters than in [RV08b]). Let $c_1 > 0$ to be taken sufficiently small depending on p_0, and let $\alpha > 0$ a sufficiently small constant to be chosen later. We have

$$\begin{aligned}
\mathbb{P}\left(\sum_{j=1}^n |\zeta_j|^2 \le c_1 \varepsilon_0^2 n \right) &= \mathbb{P}\left(n - \frac{1}{c_1 \varepsilon_0^2} \sum_{j=1}^n |\zeta_j|^2 \ge 0 \right) \\
&\le \mathbb{E} \exp\left(c_1 \alpha n - \frac{\alpha}{\varepsilon_0^2} \sum_{j=1}^n |\zeta_j|^2 \right) \\
&= e^{c_1 \alpha n} \prod_{j=1}^n \mathbb{E} \exp\left(-\alpha |\zeta_j|^2 / \varepsilon_0^2 \right). \qquad (2.10)
\end{aligned}$$

For arbitrary $j \in [n]$ we have

$$\begin{aligned}
\mathbb{E} \exp\left(-\alpha |\zeta_j|^2 / \varepsilon_0^2 \right) &= \int_0^1 \mathbb{P}\big(\exp\left(-\alpha |\zeta_j|^2 / \varepsilon_0^2 \right) \ge u \big) \mathrm{d}u \\
&= \int_0^\infty \mathbb{P}\big(|\zeta_j| \le s\varepsilon_0 / \sqrt{\alpha} \big) \mathrm{d}(e^{-s^2}) \\
&\le p_0 \int_0^{\sqrt{\alpha}} \mathrm{d}(e^{-s^2}) + \int_{\sqrt{\alpha}}^\infty \mathrm{d}(e^{-s^2}) \\
&= p_0(1 - e^{-\alpha}) + e^{-\alpha} \\
&= 1 - (1 - p_0)(1 - e^{-\alpha}).
\end{aligned}$$

Inserting this in (2.10), we obtain

$$\mathbb{P}\left(\sum_{j=1}^{n}|\zeta_j|^2 \le c_1\varepsilon_0^2 n\right) \le e^{c_1\alpha n}\big[1-(1-p_0)(1-e^{-\alpha})\big]^n \le \exp\big(n\big(c_1\alpha-(1-p_0)(1-e^{-\alpha})\big)\big).$$

The claim now follows by setting $c_1 = (1-p_0)/2$ (for instance) and taking α a sufficiently small constant. □

Let $M = A\circ X + B$ be as in Definition 1.4. In the following lemmas we assume that the atom variable ξ has κ-controlled second moment for some fixed $\kappa \ge 1$. For $v \in \mathbb{C}^m$ and $i \in [n]$ we write

$$v^i := (v_j a_{ij})_{j=1}^m \tag{2.11}$$

For $\alpha > 0$ we denote

$$I_\alpha(v) := \{i \in [n] : \|v^i\| \ge \alpha\}. \tag{2.12}$$

Lemma 2.9 (Crude anti-concentration for the image of a fixed vector)**.** *Fix $v \in \mathbb{C}^m$ and let $\alpha > 0$ such that $I_\alpha(v) \ne \varnothing$. For all $I_0 \subset I_\alpha(v)$,*

$$\sup_{w\in\mathbb{C}^n} \mathbb{P}_{I_0}\left(\|Mv-w\| \le c_0\alpha|I_0|^{1/2}\right) \le e^{-c_0|I_0|} \tag{2.13}$$

where $c_0 > 0$ is a constant depending only on κ (recall our notation $\mathbb{P}_{I_0}(\cdot)$ from Section 1.7).

Proof. Fix $w \in \mathbb{C}^n$ arbitrarily. For any $i \in I_\alpha(v)$ and any $t \ge 0$ we have

$$\mathbb{P}(|R_i\cdot v - w_i| \le t) \;\le\; p_{\xi,v^i}(t) \;=\; p_{\xi,v^i/\|v^i\|}(t/\|v^i\|) \;\le\; p_{\xi,v^i/\|v^i\|}(t/\alpha).$$

Taking $t = \alpha r_0$, by Lemma 2.6 we have

$$\mathbb{P}(|R_i\cdot v - w_i| \le \alpha r_0) \le 1 - r_0 \tag{2.14}$$

where $r_0 > 0$ depends only on κ.

Fix $I_0 \subset I_\alpha(v)$ arbitrarily. We may assume without loss of generality that I_0 is non-empty. By Lemma 2.8(a) there exists $c_1 > 0$ depending only on κ such that

$$\mathbb{P}_{I_0}\left(\sum_{i\in I_0}|R_i\cdot v - w_i|^2 \le c_1 r_0^2\alpha^2|I_0|\right) \le e^{-c_1|I_0|}. \tag{2.15}$$

Now for any $\tau \ge 0$,

$$\mathbb{P}_{I_0}\left(\|Mv-w\| \le \tau|I_0|^{1/2}\right) = \mathbb{P}_{I_0}\left(\sum_{i=1}^{n}|R_i\cdot v - w_i|^2 \le \tau^2|I_0|\right) \le \mathbb{P}_{I_0}\left(\sum_{i\in I_0}|R_i\cdot v - w_i|^2 \le \tau^2|I_0|\right)$$

and the claim follows by taking $\tau = c_1^{1/2} r_0\alpha =: c_0\alpha$ and applying (2.15). □

By similar lines, using Lemmas 2.8(b) and 2.7 in place of Lemmas 2.8(a) and 2.6, respectively, one obtains the following, which is superior to Lemma 2.9 for vectors v with small ℓ_∞ norm. The details are omitted.

Lemma 2.10 (Improved anti-concentration for the image of a fixed vector)**.** *Fix $v \in \mathbb{C}^m$. Let $\alpha > 0$ such that $I_\alpha(v) \neq \varnothing$ and fix $I_0 \subset I_\alpha(v)$ nonempty. For all $t \geq 0$,*

$$\sup_{w \in \mathbb{C}^n} \mathbb{P}_{I_0}\left(\|Mv - w\| \leq t|I_0|^{1/2} \right) \leq O_\kappa\left(\frac{1}{\alpha}\big(t + \|v\|_\infty\big) \right)^{|I_0|}. \tag{2.16}$$

3. Invertibility from connectivity: Compressible vectors

In this section we combine the anti-concentration estimates from Section 2 with union bounds over ε-nets (as obtained for instance from Lemma 2.2) to prove that with high probability, a random matrix M as in Theorem 1.13 or Theorem 1.25 is well-invertible on the set of compressible vectors $\mathrm{Comp}(\theta, \rho)$ (as defined in (2.1)) for appropriate choices of θ, ρ. Hence, there will be a competition between the quality of the anti-concentration estimates and the cardinality of the ε-nets. For small values of θ we can use ε-nets of small cardinality, but only have poor anti-concentration bounds (namely, Lemma 2.9), while for large θ the nets are very large, but we have access to the improved anti-concentration of Lemma 2.10.

In both cases we start with a crude result, Lemma 3.3, giving control for the vectors in $\mathrm{Comp}(\theta_0, \rho_0)$ for some small value of θ_0 (possibly depending on n). We then use an iterative argument argument to obtain control on $\mathrm{Comp}(\theta, \rho)$ for larger values of θ while lowering the parameter ρ. For Theorem 1.13 we want to take θ close to 1, while for Theorem 1.25 a constant order value of θ will suffice.

It turns out that that while the standard ε-net construction (see Lemma 2.2) suffices to prove Lemma 3.3, it is insufficient to obtain control on $\mathrm{Comp}(\theta, \rho)$ for the desired values of θ. For the broadly connected case this is essentially due to working in $\mathbb{C}^n$ rather than $\mathbb{R}^n$, which causes a factor 2 increase in metric entropies (this difficulty was not present in the proof of Theorem 1.11 in [RZ16] as they worked in $\mathbb{R}^n$). The situation is worse for the case of Theorem 1.25, the main source of difficulty being that $\|B\|$ can be of arbitrary polynomial order. As a consequence, the starting point θ_0 for our iterative argument will be of size $o(1)$. This prevents us from being able to use the third condition of the super-regularity hypothesis (see Definition 1.24), which only "sees" vectors that are essentially supported on more than εn coordinates.

We deal with this by *reducing the entropy cost* of the nets over which we take union bounds. In Section 3.2 we prove Lemma 3.5 which shows, roughly speaking, that if we have already established control on vectors in $\mathrm{Comp}(\theta, \rho)$ for some θ, ρ, then we can control the vectors in $\mathrm{Comp}(\theta + \Delta, \rho')$ for some small Δ, ρ' using a *random* net of significantly smaller cardinality than the net provided by Lemma 2.2. We can then increment θ from θ_0 up to size $\gg 1$, taking steps of size Δ. For the broadly connected case we can continue and take θ as close to 1 as desired.

The entropy reduction argument for Lemma 3.5 makes use of a strong version of the well-known Restricted Invertibility Theorem due to Spielman and Srivastava – see Theorem 3.7.

We now state the main results of this section. For $K \geq 1$ we denote the *boundedness event*

$$\mathcal{B}(K) := \big\{ \|M\| \leq K\sqrt{n} \big\}. \tag{3.1}$$

With a fixed choice of K we write

$$\mathcal{E}(\theta, \rho) := \mathcal{B}(K) \wedge \big\{ \exists u \in \mathrm{Comp}(\theta, \rho) : \|Mu\| \leq \rho K\sqrt{n} \big\}. \tag{3.2}$$

Proposition 3.1 (Compressible vectors: broadly connected profile)**.** *Let $M = A \circ X + B$ be as in Definition 1.4 with $n/2 \leq m \leq 2n$, and assume that ξ has κ-controlled second moment for some $\kappa \geq 1$ (see Definition 2.4). Let $K \geq 1$ and $\sigma_0, \delta, \nu \in (0,1)$. There exist $\theta_0(\kappa, \sigma_0, \delta, K) > 0$ and $\rho(\kappa, \sigma_0, \delta, \nu, K) > 0$ such that the following holds. Assume*

(1) $|\mathcal{N}_{A(\sigma_0)^\mathsf{T}}(j)| \geq \delta n$ for all $j \in [m]$;

(2) $|\mathcal{N}^{(\delta)}_{A(\sigma_0)^\mathsf{T}}(J)| \geq \min((1+\nu)|J|, n)$ for all $J \subset [m]$ with $|J| \geq \theta_0 m$.

Then for any $0 < \theta \leq (1 - \frac{\delta}{4}) \min(\frac{n}{m}, 1)$,

$$\mathbb{P}(\mathcal{E}(\theta, \rho)) \ll \exp\big(-c_\kappa \delta \sigma_0^2 n\big) \tag{3.3}$$

where $c_\kappa > 0$ depends only on κ and the implied constant depends only on $\kappa, \sigma_0, \delta, \nu$ and K.

The following gives control of compressible vectors for more general profiles than in Proposition 3.1 (essentially removing the condition (2)). However, we have to take the parameter ρ much smaller, and we only cover vectors that are essentially supported on a small (linear) proportion of the coordinates, rather than a proportion close to one.

Proposition 3.2 (Compressible vectors: general profile with large perturbation)**.** *Let $M = A \circ X + B$ be as in Definition 1.4 with $n/2 \leq m \leq 2n$. Assume ξ has κ-controlled second moment for some $\kappa \geq 1$, and that for some $a_0 > 0$ we have*

$$\sum_{i=1}^{n} a_{ij}^2 \geq a_0^2 n \quad \text{for all } j \in [m]. \tag{3.4}$$

Fix $\gamma \geq 1/2$ and let $1 \leq K = O(n^{\gamma - 1/2})$. Then for some $\rho = \rho(\gamma, a_0, \kappa, n) \gg_{\gamma, a_0, \kappa} n^{-O(\gamma^2)}$ and a sufficiently small constant $c_0 > 0$ we have

$$\mathbb{P}\big(\mathcal{E}(c_0 a_0^2, \rho)\big) \ll_{\gamma, a_0, \kappa} \exp\big(-c_\kappa a_0^2 n\big) \tag{3.5}$$

where $c_\kappa > 0$ depends only on κ.

3.1. **Highly compressible vectors.** In this subsection we establish the following crude version of Proposition 3.2, giving control on vectors in $\mathrm{Comp}(\theta_0, \rho_0)$ with θ_0 sufficiently small depending on a_0 and K.

Lemma 3.3 (Highly compressible vectors)**.** *Let $M = A \circ X + B$ be as in Definition 1.4 with $m \le 2n$. Assume that ξ has κ-controlled second moment for some $\kappa \ge 1$. Suppose also that there is a constant $a_0 > 0$ such that for all $j \in [m]$, $\sum_{i=1}^n a_{ij}^2 \ge a_0^2 n$. Let $K \ge 1$. Then with notation as in (3.2) we have*

$$\mathbb{P}(\mathcal{E}(\theta_0, \rho_0)) \le e^{-c_\kappa a_0^2 n} \tag{3.6}$$

where $\theta_0 = c_\kappa a_0^2 / \log(K/a_0^2)$ and $\rho_0 = c_\kappa a_0^2 / K$ for a sufficiently small $c_\kappa > 0$ depending only on κ.

We will need the following lemma, which ensures that the set $I_\alpha(v)$ from (2.12) is reasonably large when the columns of A have large ℓ_2 norm. A similar argument has been used in [LR12] and [RZ16].

Lemma 3.4 (Many good rows)**.** *Let A be an $n \times m$ matrix as in Definition 1.4, and assume that for some $a_0 > 0$ we have $\sum_{i=1}^n a_{ij}^2 \ge a_0^2 n$ for all $j \in [m]$. Then for any $v \in S^{m-1}$ we have $|I_{a_0/2}(v)| \ge \frac{1}{2} a_0^2 n$.*

Proof. Writing $\alpha = a_0/\sqrt{2}$, we have

$$\begin{aligned} a_0^2 n &\le \sum_{i=1}^n \sum_{j=1}^m |v_j|^2 a_{ij}^2 \\ &= \sum_{i \in I_\alpha(v)} \sum_{j=1}^m |v_j|^2 a_{ij}^2 + \sum_{i \notin I_\alpha(v)} \sum_{j=1}^m |v_j|^2 a_{ij}^2 \\ &\le \sum_{i \in I_\alpha(v)} \sum_{j=1}^m |v_j|^2 + \sum_{i \notin I_\alpha(v)} \frac{1}{2} a_0^2 \\ &\le |I_\alpha(v)| + \frac{1}{2} a_0^2 n \end{aligned}$$

and rearranging gives the claim. □

Proof of Lemma 3.3. Fix $J \subset [m]$ of size $\lfloor \theta_0 m \rfloor$ and let $v \in S^J$ be arbitrary. Writing $\alpha = a_0/\sqrt{2}$, by Lemma 2.9 and our choice of ρ_0 (with $c_\kappa > 0$ sufficiently small depending on κ),

$$\mathbb{P}\big(\|Mv\| \le \rho_0 K \sqrt{n}\big) \le \mathbb{P}\Big(\|Mv\| \le c_\kappa a_0 |I_\alpha(v)|^{1/2}\Big) \le e^{-c_\kappa |I_\alpha(v)|}.$$

Applying Lemma 3.4, we obtain

$$\mathbb{P}\big(\|Mv\| \le \rho_0 K \sqrt{n}\big) \le e^{-c_\kappa a_0^2 n} \qquad \forall v \in S^J \tag{3.7}$$

(adjusting c_κ). By Lemma 2.2 we may fix $\Sigma_J \subset S^J$ a $\rho_0/4$-net for S^J such that $|\Sigma_J| \le (12/\rho_0)^{2k}$. Suppose that $\|M\| \le K\sqrt{n}$ and that $\|Mu\| \le \rho_0 K \sqrt{n}$ for some

$u \in S^{m-1} \cap (\mathbb{C}^J)_{\rho_0/4}$. Let $u' \in \mathbb{C}^J$ with $\|u - u'\| \le \rho_0/4$, and let $u'' \in \Sigma_J$ with $\|u'' - \frac{u'}{\|u'\|}\| \le \rho_0/4$. By the triangle inequality,

$$\|u - u''\| \le \|u - u'\| + \left\|u' - \frac{u'}{\|u'\|}\right\| + \left\|\frac{u'}{\|u'\|} - u''\right\| \le 3\rho_0/4$$

where the bound on the middle term follows from $|\|u'\| - 1| \le \rho_0/4$ (also by the triangle inequality). We have

$$\|Mu''\| \le \|Mu\| + \|M(u - u'')\| \le \rho_0 K\sqrt{n} + K\sqrt{n} \cdot (3\rho_0/4) \le 2\rho_0 K\sqrt{n}.$$

Applying the union bound and (3.7) (adjusting c_κ to replace ρ_0 by $2\rho_0$),

$$\begin{aligned}\mathbb{P}\big(\exists u \in S^{m-1} \cap (\mathbb{C}^J)_{\rho_0/8} : \|Mu\| \le \rho_0 K\sqrt{n}\big) \\ &\le \mathbb{P}\big(\exists u'' \in \Sigma_J : \|Mu''\| \le 2\rho_0 K\sqrt{n}\big) \\ &\le O(1/\rho_0)^{2\theta_0 m} e^{-c_\kappa a_0^2 n}\end{aligned}$$

From (2.1) and applying the union bound over all choice of $J \in \binom{[m]}{\theta_0 m}$,

$$\begin{aligned}\mathbb{P}(\mathcal{E}(\theta_0, \rho_0/4)) &\le O(1/\theta_0)^{\theta_0 m} O(1/\rho_0)^{2\theta_0 m} e^{-c_\kappa a_0^2 n} \\ &\le O\left(\frac{1}{\theta_0 \rho_0^2}\right)^{2\theta_0 n} e^{-c_\kappa a_0^2 n}\end{aligned}$$

where we used our assumption $m \le 2n$. The desired bound now follows from substituting our choices of θ_0, ρ_0, and again adjusting the constant c_κ to replace $\rho_0/4$ by ρ_0 in the above. □

3.2. **An entropy reduction lemma.** The aim of this subsection is to establish the following:

Lemma 3.5 (Control by a random net of small cardinality)**.** *For every $I \subset [n], J \subset [m]$, $\varepsilon > 0$ there is a random finite set $\Sigma_{I,J}(\varepsilon) \subset S^J$, measurable with respect to $\mathcal{F}_{I.J} = \langle \{\xi_{ij}\}_{i \in I, j \in J} \rangle$, such that the following holds. Let $\rho \in (0,1)$, $K > 0$ and $0 < \theta < \frac{n}{m}$. On $\mathcal{B}(K) \wedge \mathcal{E}(\theta, \rho)^c$, for all $J \subset [m]$ with $|J| > \theta m$ and all $\beta, \rho' \in (0,1)$, putting*

$$\rho'' = \frac{6\rho'}{\beta\rho}\left(\frac{n}{\lfloor \theta m \rfloor}\right)^{1/2} \tag{3.8}$$

there exists $I \subset [n]$ with $|I| = \lfloor (1-\beta)^2 \lfloor \theta m \rfloor \rfloor$ such that

(1) $|\Sigma_{I,J}(\rho'')| \le (C/\rho'')^{2(|J|-|I|)}$ for an absolute constant $C > 0$, and
(2) for any $u \in S^{m-1} \cap (\mathbb{C}^J)_{\rho'}$ such that $\|Mu\| \le \rho' K\sqrt{n}$, we have $\operatorname{dist}(u, \Sigma_{I,J}(\rho'')) \le 3\rho''$.

Furthermore, writing

$$\mathcal{G}_{I,J}(\rho'') := \left\{ \left|\Sigma_{I,J}(\rho'')\right| \le \left(\frac{C}{\rho''}\right)^{2(|J|-|I|)} \right\} \tag{3.9}$$

we have that for any $\theta' \in (\theta, 1]$,

$$\mathcal{E}(\theta,\rho)^c \wedge \mathcal{E}(\theta',\rho') \subset \bigvee_{J \in \binom{[m]}{\theta' m}} \bigvee_{I \in \binom{[n]}{(1-\beta)^2 \lfloor \theta m \rfloor}} \Big(\mathcal{G}_{I,J}(\rho'') \wedge \Big\{ \exists u \in \Sigma_{I,J}(\rho'') : \|Mu\| \le 4\rho'' K\sqrt{n} \Big\} \Big). \tag{3.10}$$

Remark 3.6. We obtain the random set $\Sigma_{I,J}(\varepsilon)$ as the intersection of the sphere S^J with an ε-net of the kernel of the submatrix $M_{I,J}$. However, for our purposes it only matters that it is fixed by conditioning on the rows $\{R_i\}_{i \in I}$, has small cardinality, and serves as a net for almost-null vectors of M that are supported on J.

To prove Lemma 3.5 we use the following version of the Restricted Invertibility Theorem due to Spielman and Srivastava [SS12] (the version stated below is taken from [MSS14, Theorem 3.1]), which strengthens the classic result of Bourgain and Tzafriri [BT87].

Theorem 3.7 (Restricted Invertibility Theorem)**.** *Suppose* $v_1, \dots, v_n \in \mathbb{C}^m$ *are such that* $\sum_{i=1}^n v_i v_i^* = I_m$. *For any* $\beta \in (0,1)$, *there is a subset* $I \subset [n]$ *of size* $|I| = \lfloor (1-\beta)^2 m \rfloor$ *for which*

$$\lambda_{|I|}\Big(\sum_{i \in I} v_i v_i^* \Big) \ge \beta^2 m/n \tag{3.11}$$

where $\lambda_k(A)$ *denotes the* k*th largest eigenvalue of a Hermitian matrix* A.

This has the following consequence, which can be seen as a robust quantitative version of the basic fact from linear algebra that the row rank of a matrix is equal to its column rank.

Corollary 3.8. *Let* M *be an* $n \times m$ *matrix with* $n \ge m$, *and assume* $s_m(M) \ge \varepsilon_0 \sqrt{n}$ *for some* $\varepsilon_0 > 0$. *For any* $\beta \in (0,1)$ *there exists* $I \subset [n]$ *with* $|I| = \lfloor (1-\beta)^2 m \rfloor$ *such that*

$$s_{|I|}(M_{I,[m]}) \ge \beta \varepsilon_0 \sqrt{m}.$$

Remark 3.9. The original Restricted Invertibility Theorem of Bourgain and Tzafriri [BT87] only gives $|I| \ge cm$ and $s_{|I|}(M_{I,[m]}) \ge c\varepsilon_0\sqrt{m}$ for some (small) absolute constant $c > 0$, while it will be important for our purposes to be able to take I of size close to m.

Proof of Corollary 3.8. By the singular value decomposition it suffices to consider M of the form $M = U\Sigma$ where U is an $n \times m$ matrix with orthonormal columns and Σ is an $m \times m$ diagonal matrix with entries bounded below by $\varepsilon_0\sqrt{n}$. Fix $\alpha \in (0,1)$. Letting $v_1^*, \dots, v_n^* \in \mathbb{C}^m$ denote the rows of U, it follows from orthonormality that

$$I_m = U^*U = \sum_{i=1}^n v_i v_i^*.$$

Hence, we can apply Theorem 3.7 to obtain a subset $I \subset [n]$ with $|I| = \lfloor (1-\beta)^2 m \rfloor$ such that

$$s_{|I|}(U_{I,[m]})^2 = \lambda_{|I|}\Big(\sum_{i\in I} v_i v_i^*\Big) \geq \beta^2 m/n.$$

Now we have

$$s_{|I|}(M_{I,m}) \geq s_{|I|}(U_{I,m}) s_m(\Sigma) \geq \beta\sqrt{\frac{m}{n}}\varepsilon_0\sqrt{n} = \beta\varepsilon_0\sqrt{m}. \qquad \square$$

Proof of Lemma 3.5. Let $I \subset [n], J \subset [m]$, and write $V_{I,J} = \mathbb{C}^J \cap \ker(M_{I,J})$. Conditional on $\mathcal{F}_{I,J}$, for $\varepsilon > 0$ we let $\Sigma_{I,J}(\varepsilon)$ be an ε-net of $S^{m-1} \cap V_{I,J}$. By Lemma 2.2 we may take

$$|\Sigma_{I,J}(\varepsilon)| = O(1/\varepsilon)^{2\dim(V_{I,J})}. \tag{3.12}$$

Let $\rho, \rho' \in (0,1)$, $K > 0$ and $0 < \theta < \frac{n}{m}$. Fix $\beta \in (0,1)$ and $J \subset [m]$ with $|J| > \theta m$. On $\mathcal{E}(\theta,\rho)^c$, for all $J_0 \subset J$ with $|J_0| = \lfloor \theta m \rfloor$ we have

$$s_{\lfloor \theta m \rfloor}(M_{[n],J_0}) \geq \rho K\sqrt{n}.$$

By Corollary 3.8 there exists $I \subset [n]$ with $|I| = \lfloor (1-\beta)^2 \lfloor \theta m \rfloor \rfloor$ such that

$$s_{|I|}(M_{I,J_0}) \geq \beta\rho K\sqrt{\lfloor \theta m \rfloor}.$$

By the Cauchy interlacing law,

$$s_{|I|}(M_{I,J}) \geq \beta\rho K\sqrt{\lfloor \theta m \rfloor}. \tag{3.13}$$

In particular, the submatrix $(y_{ij})_{i\in I, j\in J}$ has full row-rank, which implies $\dim(V_{I,J}) = |J| - |I|$. From (3.12) we conclude

$$|\Sigma_{I,J}(\varepsilon)| = O(1/\varepsilon)^{2(|J|-|I|)} \tag{3.14}$$

for any $\varepsilon > 0$.

Now suppose there exists $u \in S^{m-1} \cap (\mathbb{C}^J)_{\rho'}$ such that

$$\|Mu\| \leq \rho' K\sqrt{n}. \tag{3.15}$$

Letting $v' \in \mathbb{C}^J$ such that $\|u - v'\| \leq \rho'$, and putting $v := v'/\|v'\| \in S^J$, by the triangle inequality we have $\|u - v\| \leq 2\rho'$ and

$$\|Mv\| \leq \|Mu\| + \|M\|\|u - v\| \leq 3\rho' K\sqrt{n}. \tag{3.16}$$

On the other hand,

$$\|Mv\| \geq \|M_{I,[m]}v\| = \|M_{I,[m]}(\mathrm{I} - P_{V_{I,J}})v\|$$

where $P_{V_{I,J}}$ is the matrix for orthogonal projection to the subspace $V_{I,J}$. Applying (3.13),

$$\|Mv\| \geq \|(\mathrm{I} - P_{V_{I,J}})v\|\beta\rho K\sqrt{\lfloor \theta m \rfloor}.$$

Together with (3.16) this implies that v lies within distance

$$\frac{3\rho'\sqrt{n}}{\beta\rho\sqrt{\lfloor \theta m \rfloor}} = \rho''/2 \tag{3.17}$$

of the subspace $V_{I,J}$. Since v is a unit vector we have $\operatorname{dist}(v, S^{m-1} \cap V_{I,J}) \le \rho''$ by the triangle inequality, and

$$\begin{aligned}\operatorname{dist}(u, \Sigma_{I,J}(\rho'')) &\le \|u - v\| + \rho'' + \operatorname{dist}(v, S^{m-1} \cap V_{I,J}) \\ &\le 2\rho' + 2\rho'' \le 3\rho''\end{aligned}$$

as desired (that $2\rho' \le \rho''$ follows from inspection of (3.8)).

Now to prove (3.10), let $\theta' \in (\theta, 1]$. Intersecting with $\mathcal{E}(\theta,\rho)^c$ and applying the first part of the lemma,

$$\begin{aligned}&\mathcal{E}(\theta,\rho)^c \wedge \mathcal{E}(\theta',\rho') \\ &\quad = \mathcal{B}(K) \wedge \mathcal{E}(\theta,\rho)^c \wedge \bigvee_{J \in \binom{[m]}{\theta' m}} \Big\{ \exists v \in (S^J)_{\rho'} : \|Mv\| \le \rho' K \sqrt{n} \Big\} \\ &\quad \subset \bigvee_{J \in \binom{[m]}{\theta' m}} \bigvee_{I \in \binom{[n]}{(1-\beta)^2 \lfloor \theta m \rfloor}} \Big(\mathcal{G}_{I,J}(\rho'') \wedge \Big\{ \exists u \in \Sigma_{I,J}(\rho'') : \|Mu\| \le 4\rho'' K \sqrt{n} \Big\} \Big)\end{aligned} \tag{3.18}$$

where in the last line we noted that for $v \in (S^J)_{\rho'}, u \in \Sigma_{I,J}(\rho'')$ such that $\|u - v\| \le 3\rho''$, we have

$$\|Mu\| \le \|Mv\| + 3\rho'' K\sqrt{n} \le (\rho' + 3\rho'')K\sqrt{n} \le 4\rho'' K \sqrt{n}. \qquad \square$$

3.3. **Broadly connected profile: Proof of Proposition 3.1.** We will obtain Proposition 3.1 from an iterative application of the following lemma:

Lemma 3.10 (Incrementing compressibility: broadly connected profile)**.** *Let $M = A \circ X + B$ be as in Definition 1.4 with $m \ge n/2$. Assume ξ has κ-controlled second moment for some $\kappa \ge 1$, and that for some $\sigma_0, \delta, \nu, \theta_1 \in (0,1)$ we have*

(1) $|\mathcal{N}_{A(\sigma_0)}(j)| \ge \delta n$ for all $j \in [m]$;
(2) $|\mathcal{N}^{(\delta)}_{A(\sigma_0)}(J)| \ge \min((1+\nu)|J|, n)$ for all $J \subset [m]$ with $|J| \ge (\theta_1/2)m$.

Let $K \ge 1$, $\rho \in (0,1)$, and $\theta \in [\theta_1, 1)$ such that $(1 + \frac{\nu}{2})\theta m < n$. There exists $\rho' = \rho'(\kappa, \sigma_0, \delta, \nu, \rho, \theta, K) > 0$ such that

$$\mathbb{P}\Big(\mathcal{E}(\theta,\rho)^c \wedge \mathcal{E}\Big(\Big(1 + \frac{\nu}{10}\Big)\theta, \rho'\Big)\Big) = O(e^{-n}) \tag{3.19}$$

where the implied constant depends only on $\kappa, \sigma_0, \delta, \nu, \rho, \theta$, and K.

Proof. We may assume n is sufficiently large depending on $\kappa, \sigma_0, \delta, \nu, \rho, \theta, K$. Write $\theta' = \big(1 + \frac{\nu}{10}\big)\theta$ and take $\beta = \frac{\nu}{10}$. Let $\rho' > 0$ to be taken sufficiently small depending on $\kappa, \sigma_0, \delta, \nu, \rho, \theta, K$, and let ρ'' be as in (3.8). Intersecting the right hand side of

(3.10) with $\mathcal{E}(\theta,\rho)^c$, we have

$$\begin{aligned}
&\mathcal{E}(\theta,\rho)^c \wedge \mathcal{E}(\theta',\rho') \subset \\
&\quad \bigvee_{J\in\binom{[m]}{\theta' m}} \bigvee_{I\in\binom{[n]}{(1-\beta)^2\lfloor\theta m\rfloor}} \mathcal{G}_{I,J}(\rho'') \wedge \mathcal{E}(\theta,\rho)^c \wedge \Big\{\exists u \in \Sigma_{I,J}(\rho'') : \|Mu\| \le 4\rho'' K\sqrt{n}\Big\} \\
&\subset \bigvee_{J\in\binom{[m]}{\theta' m}} \bigvee_{I\in\binom{[n]}{(1-\beta)^2\lfloor\theta m\rfloor}} \mathcal{G}_{I,J}(\rho'') \wedge \Big\{\exists u \in \Sigma_{I,J}(\rho'') \setminus \mathrm{Comp}(\theta,\rho) : \|Mu\| \le 4\rho'' K\sqrt{n}\Big\}
\end{aligned} \tag{3.20}$$

where the second line follows by taking ρ' small enough that $4\rho'' < \rho$.

Fix $J \subset [m]$ and $I \subset [n]$ of sizes $\lfloor\theta' m\rfloor, \lfloor(1-\beta)^2\lfloor\theta m\rfloor\rfloor$, respectively, and condition on $\mathcal{F}_{I,[n]}$ (recall the notation (1.33)) to fix $\Sigma_{I,J}(\rho'')$. Consider an arbitrary element $u \in \Sigma_{I,J}(\rho'') \setminus \mathrm{Comp}(\theta,\rho)$. By Lemma 2.1, there is a set $L \subset [m]$ with $|L| \ge (1 - \frac{\nu}{C_0^2})\theta m$ and

$$\frac{\rho}{\sqrt{m}} \le |u_j| \le \frac{C_0}{\sqrt{\nu\theta m}} \tag{3.21}$$

for all $j \in L$, where $C_0 > 0$ is an absolute constant to be taken sufficiently large. For any $i \in \mathcal{N}^{(\delta)}(L)$, we have

$$\|(u_L)^i\|^2 \ge \sum_{i\in L : a_{ij}\ge\sigma_0} |u_j|^2 a_{ij}^2 \ge \frac{\rho^2}{m}\sigma_0^2\delta|L| \ge \frac{1}{2}\rho^2\sigma_0^2\delta\theta =: \alpha^2 \tag{3.22}$$

where in the last inequality we took C_0 sufficiently large. Hence,

$$|I_\alpha(u_L)| \ge |\mathcal{N}^{(\delta)}(L)| \ge \min\big(n, (1+\nu)(1-\nu/C_0^2)\theta m\big) \ge \Big(1+\frac{\nu}{2}\Big)\theta m \tag{3.23}$$

taking C_0 larger if necessary, where in the second inequality we used our assumption $\theta \ge \theta_1$, and in the third inequality we used our assumption $(1+\frac{\nu}{2})\theta m < n$.

Fix $I_0 \subset I_\alpha(u_L) \setminus I$ of size $n_0 := \lfloor(1+\frac{\nu}{2})\theta m\rfloor - |I|$. In particular,

$$\begin{aligned}
\frac{\nu}{2}\theta m \le n_0 &\le \Big(1+\frac{\nu}{2}\Big)\theta m - (1-2\beta)\theta m \\
&\le \nu\theta m
\end{aligned} \tag{3.24}$$

and

$$\begin{aligned}
n_0 + 2|I| - 2|J| &\ge \Big(1+\frac{\nu}{2}\Big)\theta m + (1-2\beta)\theta m - 2\Big(1+\frac{\nu}{10}\Big)\theta m - O(1) \\
&= \frac{1}{10}\nu\theta m - O(1).
\end{aligned} \tag{3.25}$$

by our choice of β. By Lemma 2.10,

$$\mathbb{P}_{I_0}\left(\|Mu\| \le 4\rho'' K\sqrt{n}\right) \le O_\kappa\left(\frac{1}{\alpha}\left(\frac{\rho'' K\sqrt{n}}{\sqrt{|I_0|}} + \frac{1}{\sqrt{\nu\theta m}}\right)\right)^{n_0} \le O_\kappa\left(\frac{\rho'' K}{\alpha\theta^{1/2}}\right)^{n_0} \tag{3.26}$$

where in the second inequality we applied the assumption $m \geq n/2$ and assumed that n is sufficiently large that $\rho'' \gg 1/K\sqrt{n}$ (it follows from (3.8) and our assumption that ρ' is independent of n that ρ'' is bounded below independent of n).

Suppose that $\mathcal{G}_{I,J}(\rho'')$ holds. Since the bound (3.26) is uniform in the choice of I_0, we can undo the conditioning and apply the union bound over elements of $\Sigma_{I,J}(\rho'') \setminus \mathrm{Comp}(\theta,\rho)$ to find

$$\begin{aligned}\mathbb{P}\Big(\exists u \in \Sigma_{I,J}(\rho'') \setminus \mathrm{Comp}(\theta,\rho) : \|Mu\| \leq 4\rho'' K\sqrt{n}\Big) \\ &\leq O\bigg(\frac{1}{\rho''}\bigg)^{2(|J|-|I|)} O_\kappa\bigg(\frac{\rho'' K}{\alpha\theta^{1/2}}\bigg)^{n_0} \\ &= O_\kappa\bigg(\frac{K}{\alpha\theta^{1/2}}\bigg)^{n_0} O(\rho'')^{n_0+2|I|-2|J|} \\ &= O_\kappa\bigg(\frac{K}{\alpha\theta^{1/2}}\bigg)^{\nu\theta m} O(\rho'')^{\frac{1}{10}\nu\theta m - O(1)}\end{aligned}$$

where in the last line we applied the bounds (3.24) and (3.25). Since this is uniform in I, J, we can undo the conditioning on $\mathcal{F}_{I,[n]}$ and apply (3.20) with another union bound over the choices of I, J to obtain

$$\mathbb{P}\big(\mathcal{E}(\theta,\rho)^c \wedge \mathcal{E}(\theta',\rho')\big) \leq 2^{m+n} O_\kappa\bigg(\frac{K}{\alpha\theta^{1/2}}\bigg)^{\nu\theta m} O\bigg(\frac{\rho'}{\nu\rho\theta^{1/2}}\bigg)^{\frac{1}{10}\nu\theta m - O(1)} \tag{3.27}$$

where we have substituted the definition of ρ''. The result now follows by taking ρ' sufficiently small. □

Now we conclude the proof of Proposition 3.1. From our assumptions it follows that for all $j \in [m]$ we have $\sum_{i=1}^n a_{ij}^2 \geq \delta\sigma_0^2 n$. Together with our assumption $m \leq 2n$, this means we can apply Lemma 3.3 to find that

$$\mathbb{P}(\mathcal{E}(\theta_0,\rho_0)) \leq e^{-c_\kappa \delta\sigma_0^2 n} \tag{3.28}$$

where $\theta_0 = c_\kappa \delta\sigma_0^2/\log(K/\delta\sigma_0^2)$ and $\rho_0 = c_\kappa\delta\sigma_0^2/K$.

We may assume without loss of generality that $\nu \leq \delta/2$. For $l \geq 1$ set $\theta_l = (1+\frac{\nu}{10})^l\theta_0$, and let k be the smallest l such that $\theta_l \geq \theta$. We have

$$\Big(1+\frac{\nu}{2}\Big)\theta_{k-1} m \leq \Big(1+\frac{\nu}{2}\Big)\theta m \leq \Big(1-\frac{\delta^2}{16}\Big)\min(m,n) < n.$$

In particular, $(1+\nu/10)^k\theta_0 \leq (1+\nu/10)\theta \leq 1$, so

$$k \leq \frac{\log\frac{1}{\theta_0}}{\log\big(1+\frac{\nu}{10}\big)} \ll_{\kappa,\sigma_0,\delta,\nu,K} 1. \tag{3.29}$$

Applying Lemma 3.10 inductively, we have that for every $1 \leq l \leq k$ there is $\rho_l > 0$ depending only on $\kappa, \sigma_0, \delta, \nu$ and K such that

$$\mathbb{P}(\mathcal{E}(\theta_l,\rho_l) \setminus \mathcal{E}(\theta_{l-1},\rho_{l-1})) = O_{\kappa,\sigma_0,\delta,\nu,K}(e^{-n}). \tag{3.30}$$

Together with (3.28) and the union bound,

$$\begin{aligned}\mathbb{P}(\mathcal{E}(\theta,\rho)) &\le \mathbb{P}(\mathcal{E}(\theta_0,\rho_0)) + \sum_{l=1}^{k} \mathbb{P}(\mathcal{E}(\theta_l,\rho_l) \setminus \mathcal{E}(\theta_{l-1},\rho_{l-1})) \\ &\le e^{-c_\kappa \delta \sigma_0^2 n} + O_{\kappa,\sigma_0,\delta,\nu,K}(e^{-n}) = O_{\kappa,\sigma_0,\delta,\nu,K}(e^{-c_\kappa \delta \sigma_0^2 n}).\end{aligned}$$

3.4. **General profile: Proof of Proposition 3.2.** For technical reasons (essentially due to the fact that we want to allow the operator norm to have arbitrary polynomial size) the anti-concentration argument from the previous section will not suffice here, and we will need the following substitute. Roughly speaking, while previously we argued by isolating a large set of coordinates on which the vector u is "flat" (see (3.21)), here we will need to locate a set on which u is *very flat*, only fluctuating by a constant factor. This is done by a simple dyadic decomposition of the range of u, which is responsible for the loss of a logarithmic factor in the probability bound. A similar argument will be used in Section 4.2.

Lemma 3.11 (Anti-concentration for the image of an incompressible vector)**.** *Let M be as in Proposition 3.2. Let $v \in \mathrm{Incomp}(\theta,\rho)$ for some $\theta,\rho \in (0,1)$, and fix $I_0 \subset [n]$ with $|I_0| \le \frac{1}{4}a_0^2 n$. Then for all $t \ge a_0\rho/\sqrt{m}$,*

$$\sup_{w\in\mathbb{C}^n} \mathbb{P}_{[n]\setminus I_0}\left(\|Mv-w\| \le t\sqrt{n}\right) = O_\kappa\left(\frac{t\log^{1/2}(\frac{\sqrt{m}}{\rho})}{a_0^2\rho\theta^{1/2}}\right)^{\frac{1}{4}a_0^2 n}. \tag{3.31}$$

Remark 3.12. Proceeding as in the proof of Lemma 3.10 would yield

$$\sup_{w\in\mathbb{C}^n} \mathbb{P}_{[n]\setminus I_0}\left(\|Mv-w\| \le t\sqrt{n}\right) = O_\kappa\left(\frac{t}{a_0^2\rho\theta^{1/2}}\right)^{\frac{1}{4}a_0^2 n} \quad \text{for all } \ t \ge \frac{a_0}{\sqrt{\theta m}}. \tag{3.32}$$

The ability to take t down to the scale $\sim \rho/\sqrt{m}$ will be crucial in the proof of Lemma 3.13 below.

Proof. We begin by finding a set of indices on which v varies by at most a factor of 2. For $k \ge 0$ let $L_k = \{j \in [m] : 2^{-(k+1)} < |v_j| \le 2^{-k}\}$. Since $v \in \mathrm{Incomp}(\theta,\rho)$, we have

$$|L^+| := |\{j \in [m] : |v_j| \ge \rho/\sqrt{m}\}| \ge \theta m.$$

Indeed, were this not the case then v would be within distance ρ of the vector v_{L^+} whose support is smaller than θm, implying $v \in \mathrm{Comp}(\theta,\rho)$. Thus, $L^+ \subset \bigcup_{k=0}^{\ell} L_k$ for some $\ell \ll \log(\frac{\sqrt{m}}{\rho})$. By the pigeonhole principle there exists $k^* \le \ell$ such that $L^* := L_{k^*}$ satisfies

$$|L^*| \ge \frac{\theta n}{\ell} \gg \frac{\theta m}{\log(\frac{\sqrt{m}}{\rho})}. \tag{3.33}$$

Denote $I^* := I_{\frac{a_0}{2}\|v_{L^*}\|}(v_{L^*})$. By Lemma 3.4,

$$|I^*| \geq \frac{1}{2}a_0^2 n. \tag{3.34}$$

Fix $i \in I^*$. By definition of I^*,

$$\|(v^i)_{L^*}\| \geq \frac{1}{2}a_0\|v_{L^*}\| \tag{3.35}$$

and since $|v_j| \gg \rho/\sqrt{m}$ on L^*,

$$\|v_{L^*}\| \gg \frac{\rho}{\sqrt{m}}|L^*|^{1/2}. \tag{3.36}$$

Furthermore, since $a_{ij} \leq 1$ for all $j \in [m]$ and v varies by a factor at most 2 on L^*,

$$\|(v^i)_{L^*}\|_\infty \leq \|v_{L^*}\|_\infty \leq 2\frac{\|v_{L^*}\|}{|L^*|^{1/2}}. \tag{3.37}$$

Fix $w \in \mathbb{C}^n$ arbitrarily. By Lemma 2.7 and the above estimates, for all $t \geq 0$ we have

$$\begin{aligned}
\mathbb{P}(|R_i \cdot v - w_i| \leq t) &\ll_\kappa \frac{t + \|(v^i)_{L^*}\|_\infty}{\|(v^i)_{L^*}\|} \\
&\ll \frac{1}{a_0}\left(\frac{t}{\|v_{L^*}\|} + \frac{\|(v^i)_{L^*}\|_\infty}{\|v_{L^*}\|}\right) \\
&\ll \frac{1}{a_0}\left(\frac{t}{\rho}\left(\frac{m}{|L^*|}\right)^{1/2} + \frac{1}{|L^*|^{1/2}}\right) \\
&= \frac{1}{a_0}\left(\frac{m}{|L^*|}\right)^{1/2}\left(\frac{t}{\rho} + \frac{1}{\sqrt{m}}\right).
\end{aligned}$$

By Lemma 2.8,

$$\begin{aligned}
\mathbb{P}_{I^*\setminus I_0}\left(\|Mv - w\| \leq t|I^* \setminus I_0|^{1/2}\right) &\leq \mathbb{P}_{I^*\setminus I_0}\Big(\sum_{i \in I^*\setminus I_0}|R_i \cdot v - w_i|^2 \leq t^2|I^* \setminus I_0|\Big) \\
&= O_\kappa\left(\frac{t\sqrt{m}}{a_0\rho|L^*|^{1/2}}\right)^{|I^*\setminus I_0|}
\end{aligned}$$

for all $t \geq \rho/\sqrt{m}$. Substituting the lower bounds (3.33), (3.34) on $|L^*|$ and $|I^*|$ and our assumption $|I_0| \leq \frac{1}{4}a_0^2 n$,

$$\mathbb{P}_{I^*\setminus I_0}\left(\|Mv - w\| \leq \frac{1}{2}ta_0\sqrt{n}\right) = O_\kappa\left(\frac{t\log^{1/2}(\frac{\sqrt{m}}{\rho})}{a_0\rho\theta^{1/2}}\right)^{\frac{1}{4}a_0^2 n}$$

for all $t \geq \rho/\sqrt{m}$. The result now follows by replacing t with $2t/a_0$ as undoing the conditioning on the remaining rows in $[n] \setminus I_0$. □

Now we are ready to prove the analogue of Lemma 3.10 for general profiles. Whereas in the broadly connected case we obtained control on vectors in $\mathrm{Comp}((1+\beta)\theta, \rho')$ after restricting to the event that we have control on $\mathrm{Comp}(\theta, \rho)$, for small $\beta > 0$, here we will also need to assume control on $\mathrm{Comp}(\theta_0, \rho_0)$ for a fixed small θ_0 at each step. The control on $\mathrm{Comp}(\theta, \rho)$ will be used to obtain a net of low cardinality using Lemma 3.5, while the control on $\mathrm{Comp}(\theta_0, \rho_0)$ will be used to obtain good anti-concentration estimates using Lemma 3.11. (In the broadly connected case the control on $\mathrm{Comp}(\theta, \rho)$ was sufficient for both purposes.)

Lemma 3.13 (Incrementing compressibility: general profile)**.** *Let M be as in Proposition 3.2, fix $\gamma > 1/2$ and put $K = n^{\gamma - 1/2}$. Let θ_0, ρ_0 be as in Lemma 3.3, and fix $\theta \in [\theta_0, c_0 a_0^2]$, where c_0 is a sufficiently small constant (we may assume the constant c in Lemma 3.3 is sufficiently small so that this interval is non-empty). We have*

$$\mathbb{P}\Big(\mathcal{E}(\theta_0, \rho_0)^c \wedge \mathcal{E}(\theta, \rho)^c \wedge \mathcal{E}(\theta + \beta a_0^2, \rho')\Big) = O_{\gamma, a_0, \kappa}(e^{-n}) \tag{3.38}$$

for some $\rho' \gg_{\gamma, a_0, \kappa} n^{-O(\gamma)}\rho$, where we set

$$\beta = c_1 \min\left(1, \frac{1}{\gamma - 1/2}\right) \tag{3.39}$$

for a sufficiently small constant $c_1 > 0$.

Proof. Let $\rho' > 0$ to be taken sufficiently small, and let ρ'' be as in (3.8). We denote $\theta' = \theta + \beta a_0^2$. Intersecting both sides of (3.10) with $\mathcal{E}(\theta_0, \rho_0)^c$, we have

$$\mathcal{E}(\theta_0, \rho_0)^c \wedge \mathcal{E}(\theta, \rho)^c \wedge \mathcal{E}(\theta', \rho') \subset$$
$$\bigvee_{J \in \binom{[m]}{\theta' m}} \bigvee_{I \in \binom{[n]}{(1-\beta)^2 \lfloor \theta m \rfloor}} \mathcal{G}_{I,J}(\rho'') \wedge \Big\{\exists u \in \Sigma_{I,J}(\rho'') \setminus \mathrm{Comp}(\theta_0, \rho_0) : \|Mu\| \le 4\rho'' K\sqrt{n}\Big\} \tag{3.40}$$

where we have assumed ρ' is small enough that $4\rho'' < \rho_0$.

Fix $J \subset [m]$ and $I \subset [n]$ of size $\lfloor \theta' m \rfloor$, $\lfloor (1-\beta)^2 \lfloor \theta m \rfloor \rfloor$, respectively, and condition on $\mathcal{F}_{I,[n]}$ to fix $\Sigma_{I,J}(\rho'')$. Fix an arbitrary $u \in \Sigma_{I,J}(\rho'') \setminus \mathrm{Comp}(\theta_0, \rho_0)$. From Lemma 3.11 we have

$$\mathbb{P}_{[n]\setminus I}\left(\|Mu\| \le 4\rho'' K\sqrt{n}\right) = O_\kappa\left(\frac{\rho'' K \log^{1/2}(\frac{\sqrt{n}}{\rho_0})}{a_0^2 \rho_0 \theta_0^{1/2}}\right)^{\frac{1}{4}a_0^2 n} \tag{3.41}$$

provided

$$\rho'' \ge \frac{c a_0 \rho_0}{K\sqrt{n}} \tag{3.42}$$

for some small constant $c > 0$ (note that we used our assumption $n/2 \le m \le 2n$).

Applying the union bound over the choices of $u \in \Sigma_{I,J}(\rho'') \setminus \mathrm{Comp}(\theta_0, \rho_0)$, on the event $\mathcal{G}_{I,J}(\rho'')$ we have

$$\begin{aligned}\mathbb{P}\left(\exists u \in \Sigma_{I,J}(\rho'') \setminus \mathrm{Comp}(\theta_0,\rho_0) : \|Mu\| \le 4\rho'' K\sqrt{n}\right) \\ &\le O\left(\frac{1}{\rho''}\right)^{2(|J|-|I|)} O_\kappa\left(\frac{\rho'' K \log^{1/2}(\frac{\sqrt{n}}{\rho_0})}{a_0^2\rho_0\theta_0^{1/2}}\right)^{\frac{1}{4}a_0^2 n} \\ &= O\left(\frac{1}{\rho''}\right)^{2(|J|-|I|)} O_{\kappa,a_0}\left(\rho'' K^2 \log(K\sqrt{n})\right)^{\frac{1}{4}a_0^2 n}\end{aligned}$$

where in the second line we substituted the expressions for ρ_0, θ_0 from Lemma 3.3. Denoting $\varepsilon = \rho'' K^2$, the above bound rearranges to

$$O_{\kappa,a_0}(\log n)^n n^{O(\gamma)} n^{O(\gamma-1/2)(|J|-|I|)} \varepsilon^{\frac{1}{4}a_0^2 n - 2(|J|-|I|)}. \tag{3.43}$$

We can bound

$$\begin{aligned}|J| - |I| &= \theta m + \beta a_0^2 m - (1-\beta)^2\theta m + O(1) \\ &\le \beta a_0^2 m + 2\beta\theta m + O(1) \\ &= O(\beta a_0^2 m) + O(1)\end{aligned}$$

where we used our assumption that $\theta \le c_0 a_0^2$. In particular, $|J|-|I| \le \frac{1}{8}a_0^2 n + O(1)$ if the constant c_1 in (3.39) is sufficiently small, and (3.43) is bounded by

$$O_{\kappa,a_0}(\log n)^n n^{O(\gamma)} n^{O(\gamma-1/2)\beta a_0^2 m} \varepsilon^{\frac{1}{8}a_0^2 n - O(1)}. \tag{3.44}$$

Applying the union bound over the choices of I, J in (3.40), which incurs a harmless factor of $2^{m+n} = O(1)^n$, and substituting the expression (3.39) for β we have

$$\mathbb{P}\Big(\mathcal{E}(\theta_0,\rho_0)^c \wedge \mathcal{E}(\theta,\rho)^c \wedge \mathcal{E}(\theta+\beta a_0^2, \rho')\Big) = O_{\kappa,a_0}(\log n)^n n^{O(\gamma)} \varepsilon^{-O(1)} (n^{O(c_1)}\varepsilon^{1/8})^{a_0^2 n}. \tag{3.45}$$

It only remains to check that we can take ε sufficiently small to obtain (3.38). From (3.42) we are constrained to take

$$\varepsilon = \rho'' K^2 \ge \frac{c a_0 \rho_0 K}{\sqrt{n}} = \frac{c' a_0^3}{\sqrt{n}}$$

for some constant $c' \in (0,1)$ sufficiently small. Taking $\varepsilon = a_0^3/\sqrt{n}$ and c_1 sufficiently small we have

$$\mathbb{P}\Big(\mathcal{E}(\theta_0,\rho_0)^c \wedge \mathcal{E}(\theta,\rho)^c \wedge \mathcal{E}(\theta+\beta a_0^2, \rho')\Big) \le O_{\kappa,a_0}(1)^n n^{O(\gamma)} n^{-.01 a_0^2 n} \tag{3.46}$$

which yields (3.38) as desired. With this choice of ε,

$$\rho' \gg \rho''\beta\rho\theta \ge \rho''\beta\rho\theta_0 \gg_{\kappa,a_0,\gamma} \rho n^{-2\gamma+1/2-o(1)}$$

as desired (recall that $\theta_0 \gg_\kappa a_0^2/\log(K/a_0) \gg_{\gamma,a_0,\kappa} 1/\log n$). □

Now we conclude the proof of Proposition 3.2. Since the event $\mathcal{B}(K)$ is monotone under increasing K, by perturbing γ and assuming n is sufficiently large we may take $K = n^{\gamma-1/2}$ with $\gamma > 1/2$. Let ρ_0, θ_0 be as in Lemma 3.3, and for $l \geq 1$ we let $\theta_l = \theta_0 + l\beta a_0^2$ with $\beta = \beta(\gamma)$ as in (3.39). By Lemma 3.13 we can inductively define a sequence ρ_l such that for each $l \geq 1$ such that $\theta_l \leq c_0 a_0^2$,

$$\rho_l \gg_{\gamma,a_0,\kappa} n^{-O(\gamma)}\rho_{l-1}$$

and

$$\mathbb{P}(\mathcal{E}(\theta_0,\rho_0)^c \wedge \mathcal{E}(\theta_{l-1},\rho_{l-1})^c \wedge \mathcal{E}(\theta_l,\rho_l)) = O_{\gamma,a_0,\kappa}(e^{-n}).$$

Applying the union bound, for some $k = O(\gamma)$ we have

$$\begin{aligned}\mathbb{P}\big(\mathcal{E}(c_0a_0^2,\rho)\big) &\leq \mathbb{P}(\mathcal{E}(\theta_0,\rho_0)) + \sum_{l=1}^{k}\mathbb{P}(\mathcal{E}(\theta_0,\rho_0)^c \wedge \mathcal{E}(\theta_{l-1},\rho_{l-1})^c \wedge \mathcal{E}(\theta_l,\rho_l))\\ &\leq e^{-c_\kappa a_0^2 n} + O_{\gamma,a_0,\kappa}(e^{-n})\\ &= O_{\gamma,a_0,\kappa}(e^{-c_\kappa a_0^2 n})\end{aligned}$$

and

$$\rho \gg_{\gamma,a_0,\kappa} n^{-O(\gamma^2)}.$$

This concludes the proof of Proposition 3.2.

4. Invertibility from connectivity: Incompressible vectors

In this section we conclude the proofs of Theorems 1.13 and 1.25 by bounding the event that $\|Mu\|$ is small for some incompressible vector u (recall the terminology from Section 2.1). We follow the (by now standard) approach of reducing to the event that a fixed row R_i of M lies close to the span of the remaining rows, an idea which goes back to the work of Komlós on the singularity probability for Bernoulli matrices [Kom67, Kom68, Kom]. This can in turn be controlled by the event that a random walk $R_i \cdot v$ concentrates near a particular point, where v is a fixed unit vector in the orthocomplement of the remaining rows. Independence of the rows allows us to condition on v, and our results from the previous section allow us to argue that v is incompressible.

For the case that the entries of R_i have variances uniformly bounded below, we could then complete the proof by applying the anti-concentration estimate of Lemma 2.7. In the present setting, however, a proportion $1-\delta$ of the entries of R_i may have zero variance. For the case of broadly connected profile we follow the argument of Rudelson and Zeitouni [RZ16] and use Proposition 3.1 to show v has essential support of size $(1-\delta/2)n$, and hence has non-trivial overlap with the support of R_i.

For the case of a super-regular profile, Proposition 3.2 only gives that v has essential support of size $\gg \delta\sigma_0^2$. In Lemma 4.1 we make use of a double counting argument to show that if we choose the row R_i at random, on average it will have good overlap with the corresponding normal vector $v^{(i)}$ (which also depends on

i). Here is where we make crucial use of the super-regularity hypothesis on A. Lemma 4.1 is a natural extension of a double counting argument used by Komlós in his work on the singularity probability for Bernoulli matrices, and which was applied to bound the smallest singular value of iid matrices by Rudelson and Vershynin in [RV08b]. We were also inspired by a similar refinement of the double counting argument from the recent paper [LLT+17] on the singularity probability for adjacency matrices of random regular digraphs.

4.1. **Proof of Theorem 1.13.** By Lemma 2.5 and multiplying X and B by a phase (which does not affect our hypotheses) we may assume that ξ has $O(\kappa_0)$-controlled second moment. Fix $K \geq 1$, and let $\rho = \rho(\kappa, \sigma_0, \delta, \nu, K)$ be as in Proposition 3.1. We may assume n is sufficiently large depending on $\kappa, \sigma_0, \delta, \nu, K$. For the remainder of the proof we restrict to the event $\mathcal{B}(K) = \{\|M\| \leq K\sqrt{n}\}$.

For $j \in [n]$ let $M^{(i)}$ denote the $n-1 \times n$ matrix obtained by removing the ith row from M. Define the good event

$$\mathcal{G} = \Big\{\forall i \in [n], \forall u \in \mathrm{Comp}(1-\delta/2, \rho),\ \|u^* M\|, \|M^{(i)} u\| > \rho K \sqrt{n}\Big\}. \tag{4.1}$$

Applying Proposition 3.1 to M^* and $M^{(i)}$ for each $i \in [n]$ (using our restriction to $\mathcal{B}(K)$) and the union bound we have

$$\mathbb{P}(\mathcal{G}) = 1 - O_{\kappa,\sigma_0,\delta,\nu,K}(ne^{-c_\kappa \delta \sigma_0^2 n}) = 1 - O_{\sigma_0,\delta,\nu,K}(e^{-c_\kappa \delta \sigma_0^2 n}) \tag{4.2}$$

adjusting c_κ slightly. Let $t \leq 1$, and define the event

$$\mathcal{E}(t) = \mathcal{G} \wedge \big\{\exists u \in \mathrm{Incomp}(1/10, \rho) : \|u^* M\| \leq t/\sqrt{n}\big\}. \tag{4.3}$$

For n sufficiently large (larger than $1/\rho K$) it suffices to show

$$\mathbb{P}(\mathcal{E}(t)) \ll_{\kappa,\sigma_0,\delta,\nu,K} t + n^{-1/2}. \tag{4.4}$$

Recalling that R_i denotes the ith row of M, we denote

$$R_{-i} = \mathrm{span}(R_j : j \in [n] \setminus \{i\}) \tag{4.5}$$

and let

$$\mathcal{E}_i(t) = \mathcal{G} \wedge \{\mathrm{dist}(R_i, R_{-i}) \leq t/\rho\}. \tag{4.6}$$

We now use a double counting argument of Rudelson and Vershynin from [RV08b] to control $\mathcal{E}(t)$ in terms of the events $\mathcal{E}_i(t)$. Suppose that $\mathcal{E}(t)$ holds, and let $u \in \mathrm{Incomp}(1/10, \rho)$ such that $\|u^* M\| \leq t/\sqrt{n}$. Then we must have $|u_i| \geq \rho/\sqrt{n}$ for at least $n/10$ elements $i \in [n]$. For each such i we have

$$\frac{t}{\sqrt{n}} \geq \|u^* M\| = \Big\| \sum_{j=1}^n \overline{u_j} R_j \Big\| \geq \Big\| P_{R_{-i}^\perp} \sum_{j=1}^n \overline{u_j} R_j \Big\| = |u_i| \left\| P_{R_{-i}^\perp} R_i \right\| \geq \frac{\rho}{\sqrt{n}} \mathrm{dist}(R_i, R_{-i})$$

where we denote by P_W the orthogonal projection to a subspace W. Thus, on $\mathcal{E}(t)$ we have that $\mathcal{E}_i(t)$ holds for at least $n/10$ values of $i \in [n]$, so by double

counting,

$$\mathbb{P}(\mathcal{E}(t)) \leq \frac{10}{n}\sum_{i=1}^{n}\mathbb{P}(\mathcal{E}_i(t)). \tag{4.7}$$

Now it suffices to show that for arbitrary fixed $i \in [n]$,

$$\mathbb{P}(\mathcal{E}_i(t)) \ll_{\kappa,\sigma_0,\delta,\nu,K} t + n^{-1/2}. \tag{4.8}$$

Fix $i \in [n]$ and condition on $\{R_j : j \in [n] \setminus \{i\}\}$. Draw a unit vector $u \in R_{-i}^{\perp}$ independent of R_i, according to Haar measure (say). Since $\operatorname{dist}(R_i, R_{-i}) \leq |R_i \cdot u|$, it suffices to show

$$\mathbb{P}(|R_i \cdot u| \leq t/\rho) \ll_{\kappa,\sigma_0,\delta,\nu,K} t + n^{-1/2}. \tag{4.9}$$

Since $u \in \ker(M^{(i)})$, on $\mathcal{G}$ we have that $u \in \operatorname{Incomp}(1-\frac{\delta}{2},\rho)$. By Lemma 2.1 there exists $L \subset [n]$ of size $|L| \geq (1-\frac{3}{4}\delta)n$ such that

$$\frac{\rho}{\sqrt{n}} \leq |u_j| \leq \frac{10}{\sqrt{\delta n}}$$

for all $j \in L$. By assumption we have $|\mathcal{N}_{A(\sigma_0)}(i)| = |\{j \in [n] : a_{ij} \geq \sigma_0\}| \geq \delta n$, so letting $J = \mathcal{N}_{A(\sigma_0)}(i) \cap L$ we have $|J| \geq \delta n/4$. Denoting $v = (u^i)_J = (a_{ij}u_j 1_{j\in J})_j$, we have

$$\|v\|^2 = \sum_{j\in J} a_{ij}^2 |u_j|^2 \geq |J|\sigma_0^2\rho^2/n \geq \delta\sigma_0^2\rho^2/4$$

and

$$\|v\|_\infty \leq \|u_J\|_\infty \leq \frac{10}{\sqrt{\delta n}}$$

(recall that $a_{ij} \leq 1$ for all $i,j \in [n]$). Conditioning on u and $\{\xi_{ij}\}_{j\notin J}$, we apply Lemma 2.7 to conclude

$$\mathbb{P}(|R_i \cdot u| \leq t/\rho) \ll_\kappa \frac{1}{\|v\|}\left(\frac{t}{\rho} + \|v\|_\infty\right) \ll \frac{1}{\rho\sigma_0\delta^{1/2}}\left(\frac{t}{\rho} + \frac{1}{\sqrt{\delta n}}\right)$$

which gives (4.9) as desired.

4.2. **Proof of Theorem 1.25.** By Lemma 2.5 and multiplying X and B by a phase (which does not affect our hypotheses) we may assume that ξ has $\kappa = O(\kappa_0)$-controlled second moment. Fix $\gamma \geq 1/2$ and let $K = O(n^{\gamma-1/2})$. We will show that for all $\tau \geq 0$,

$$\mathbb{P}\left(s_n(M) \leq \frac{\tau}{\sqrt{n}},\ \|M\| \leq K\sqrt{n}\right) \ll_{\gamma,\sigma_0,\delta,\kappa} n^{O(\gamma^2)}\tau + \sqrt{\frac{\log n}{n}}. \tag{4.10}$$

For the remainder of the proof we restrict to the boundedness event

$$\mathcal{B}(K) = \{\|M\| \leq K\sqrt{n}\}. \tag{4.11}$$

By the assumption that $A(\sigma_0)$ is (δ,ε)-super-regular we have

$$\sum_{i=1}^{n} a_{ij}^2 \geq \delta\sigma_0^2 n$$

for all $j \in [n]$. Let $a_0 = \delta^{1/2}\sigma_0$, and let $\rho = \rho(\gamma, a_0, \kappa n)$ and c_0 be as in Proposition 3.2. In particular,

$$\rho \gg_{\gamma,\delta,\sigma_0} n^{-O(\gamma^2)}. \tag{4.12}$$

Denoting $\theta = c_0\delta\sigma_0^2$, for $\tau > 0$ we define the good event

$$\mathcal{G}(\tau) = \Big\{\forall u \in \mathrm{Comp}(\theta,\rho),\ \|Mu\|, \|u^*M\| > \tau/\sqrt{n}\Big\}. \tag{4.13}$$

Applying Proposition 3.2 to M and M^*, along with the union bound, we have

$$\mathbb{P}(\mathcal{G}(\tau)) = 1 - O_{\gamma,\delta,\sigma_0,\kappa}(e^{-c_\kappa\delta\sigma_0^2 n}) \tag{4.14}$$

as long as $\tau \le \rho K n$.

Let $0 < \tau \le 1$ to be chosen later. Recalling our notation $M^{(i)}$ from Section 4.1, we define the sets

$$S_i(\tau) = \left\{u \in S^{n-1} : \|M^{(i)}u\| \le \frac{\tau}{\sqrt{n}}\right\}. \tag{4.15}$$

Informally, for small τ this is the set of unit almost-normal vectors to the subspace R_{-i} spanned by the rows of $M^{(i)}$. Writing $\mathcal{N}(i) = \mathcal{N}_{A(\sigma_0)}(i)$, we define the *good overlap events*

$$\mathcal{O}_i(\tau) = \big\{\exists u \in S_i(\tau) : |\mathcal{N}(i) \cap L^+(u,\rho)| \ge \delta\theta n\big\} \tag{4.16}$$

where

$$L^+(u) = \{j \in [n] : |u_j| \ge \rho/\sqrt{n}\}. \tag{4.17}$$

On $\mathcal{O}_i(\tau)$ we fix a vector $u^{(i)} = u^{(i)}(M^{(i)}, \tau) \in S_i(\tau)$, chosen measurably with respect to $M^{(i)}$, satisfying $|\mathcal{N}(i) \cap L^+(u,\rho)| \ge \delta\theta n$.

Lemma 4.1 (Good overlap on average)**.** *Assume $\varepsilon \le \theta/2$. Then*

$$\mathbb{P}\bigg(\mathcal{G}(\tau) \wedge \Big\{s_n(M) \le \frac{\tau}{\sqrt{n}}\Big\}\bigg) \le \frac{2}{\theta n}\sum_{i=1}^n \mathbb{P}\bigg(\mathcal{O}_i(\tau) \wedge \Big\{|R_i \cdot u^{(i)}| \le \frac{2\tau}{\rho}\Big\}\bigg). \tag{4.18}$$

Proof. Suppose $\mathcal{G}(\tau) \wedge \{s_n(M) \le \tau/\sqrt{n}\}$ holds. Then there exist $u, v \in S^{n-1}$ such that $\|Mu\|, \|M^*v\| \le \tau/\sqrt{n}$. By our restriction to $\mathcal{G}(\tau)$ we must have $u, v \in \mathrm{Incomp}(\theta,\rho)$. With notation as in (4.17) we have $|L^+(u)|, |L^+(v)| \ge \theta n$. In particular, $|L^+(u)| \ge \varepsilon n$, so

$$|\mathcal{N}(i) \cap L^+(u)| \ge \delta|L^+(u)| \ge \delta\theta n \tag{4.19}$$

for at least $(1-\varepsilon)n$ elements $i \in [n]$. Indeed, otherwise we would have

$$e_{A(\sigma_0)}(I, L^+(u)) = \sum_{i\in I}|\mathcal{N}(i) \cap L^+(u)| < \delta|I||L^+(u)|$$

for some $I \subset [n]$ with $|I| > \varepsilon n$, which contradicts our assumption that $A(\sigma_0)$ is (δ,ε)-super-regular. Since

$$\|M^{(i)}u\| \le \|Mu\| \le \frac{\tau}{\sqrt{n}}$$

for all $i \in [n]$, we have that $u \in S_i(\tau)$ for all $i \in [n]$. Thus,

$$\big|\big\{i \in L^+(v) : \mathcal{O}_i(\tau) \text{ holds}\big\}\big| \geq \theta n - \varepsilon n \geq \theta n/2. \tag{4.20}$$

Fix $i \in L^+(v)$ such that $\mathcal{O}_i(\tau)$ holds. We have

$$\frac{\tau}{\sqrt{n}} \geq \|v^* M\| \geq |v^* M u^{(i)}| \geq |v_i||R_i \cdot u^{(i)}| - \left|\sum_{j \neq i} \overline{v_j} R_j \cdot u^{(i)}\right|.$$

The first term on the right hand side is bounded below by $\frac{\rho}{\sqrt{n}}|R_i \cdot u^{(i)}|$ since $i \in L^+(v)$. By Cauchy–Schwarz the second term is bounded above by $\|M^{(i)} u^{(i)}\| \leq \tau/\sqrt{n}$, since $u^{(i)} \in S_i(\tau)$. Rearranging we conclude

$$|R_i \cdot u^{(i)}| \leq \frac{2\tau}{\rho}$$

for all $i \in L^+(v)$ such that $\mathcal{O}_i(\tau)$ holds, and the claim follows from (4.20) and double counting. □

Fix $i \in [n]$ arbitrarily, and suppose that $\mathcal{O}_i(\tau)$ holds. We condition on the rows $\{R_j\}_{j \in [n] \setminus \{i\}}$ to fix $u^{(i)}$. We begin by finding a large set on which $u^{(i)}$ is flat, following a similar dyadic pigeonholing argument as in the proof of Lemma 3.11. Letting $L_k = \{j \in [n] : 2^{-(k+1)} < |u_j^{(i)}| \leq 2^{-k}$, since

$$\delta\theta n \leq |\mathcal{N}(i) \cap L^+(u^{(i)})| \leq \left|\bigcup_{k=0}^{\ell} \mathcal{N}(i) \cap L_k\right|$$

for some $\ell \ll \log(\sqrt{n}/\rho)$, by the pigeonhole principle there exists $k^* \leq \ell$ such that $J := \mathcal{N}(i) \cap L_{k^*}$ satisfies

$$|J| \geq \delta\theta n/\ell \gg \frac{\delta\theta n}{\log(\sqrt{n}/\rho)}. \tag{4.21}$$

Let us denote $v = (a_{ij} u_j^{(i)} 1_{j \in J})_j$. Since $a_{ij} \geq \sigma_0$ for $j \in \mathcal{N}(i)$ and $|u_j^{(i)}| \gg \rho/\sqrt{n}$ for $j \in L_{k^*}$,

$$\|v\| \geq \sigma_0 \|(u^{(i)})_J\| \gg \sigma_0 \rho (|J|/n)^{1/2} \tag{4.22}$$

and since $u^{(i)}$ varies by at most a factor of 2 on J,

$$\|v\|_\infty \leq \|u^{(i)} 1_J\|_\infty \leq 2\|u^{(i)}\|/|J|^{1/2}. \tag{4.23}$$

By further conditioning on the variables $\{\xi_{ij}\}_{j \notin J}$ and applying Lemma 2.7 along with the estimates (4.22), (4.23) we have

$$\begin{aligned}\mathbb{P}\Big(|R_i \cdot u^{(i)}| \le 2\tau/\rho\Big) &\ll_\kappa \frac{\tau/\rho + \|v\|_\infty}{\|v\|} \\ &\ll \frac{1}{\sigma_0}\left(\frac{\tau/\rho}{\rho(|J|/n)^{1/2}} + \frac{1}{|J|^{1/2}}\right) \\ &= \frac{1}{\sigma_0}\left(\frac{n}{|J|}\right)^{1/2}\left(\frac{\tau}{\rho^2} + \frac{1}{\sqrt{n}}\right).\end{aligned}$$

Inserting the bound (4.21) and undoing all of the conditioning, we have shown

$$\mathbb{P}\bigg(\mathcal{O}_i(\tau) \wedge \Big\{|R_i \cdot u^{(i)}| \le \frac{2\tau}{\rho}\Big\}\bigg) \ll_\kappa \frac{1}{\sigma_0\sqrt{\delta\theta}}\left(\frac{\tau}{\rho^2} + \frac{1}{\sqrt{n}}\right)\log^{1/2}(\sqrt{n}/\rho).$$

Since the right hand side is uniform in i, applying Lemma 4.1 (taking $c_1 = c_0/2$) and substituting the expression for θ we have

$$\mathbb{P}\bigg(\mathcal{G}(\tau) \wedge \Big\{s_n(M) \le \frac{\tau}{\sqrt{n}}\Big\}\bigg) \ll_\kappa \frac{1}{\sigma_0^4\delta^2}\left(\frac{\tau}{\rho^2} + \frac{1}{\sqrt{n}}\right)\log^{1/2}(\sqrt{n}/\rho) \tag{4.24}$$

for all $\tau \ge 0$ (note that this bound is only nontrivial when $\tau \le \rho^2$, in which case our constraint $\tau \le \rho K n$ from (4.14) holds). The bound (4.10) now follows by substituting the lower bound (4.12) on ρ and the bound (4.14) on $\mathcal{G}(\tau)^c$ (which is dominated by the $O(n^{-1/2}\log^{1/2} n)$ term). This concludes the proof of Theorem 1.25.

5. Invertibility under diagonal perturbation: Proof of main theorem

In this final section we prove Theorem 1.19. See Section 1.5 for a high level discussion of the main ideas. In Sections 5.1 and 5.2 we collect the main tools of the proof: the regularity lemma, the Schur complement bound, and bounds on the operator norm of random matrices. In Section 5.3 we apply the regularity lemma to decompose the standard deviation profile A into a bounded number of submatrices enjoying various properties. In Section 5.4 we apply the decomposition to prove Theorem 1.19, on two technical lemmas, and in the final sections we prove these lemmas.

5.1. **Preliminary Tools.**

We begin by stating a version of the regularity lemma suitable for our purposes. Recall that in Theorem 1.13 we associated the standard deviation profile A with a bipartite graph. Here it will be more convenient to associate A with a directed graph. That is, to a non-negative square matrix $A = (a_{ij})_{1 \le i,j \le n}$ we associate a directed graph Γ_A on vertex set $[n]$ having an edge $i \to j$ when $a_{ij} > 0$ (note that we allow Γ_A to have self-loops, though the diagonal of A will have a negligible effect on our arguments). The notation

(1.18)–(1.19) extends to this setting. Additionally, we denote the *density* of the pair (I, J)

$$\rho_A(I, J) := \frac{e_A(I, J)}{|I||J|}.$$

Definition 5.1 (Regular pair)**.** Let A be an $n \times n$ matrix with non-negative entries. For $\varepsilon > 0$, we say that a pair of vertex subsets $I, J \subset [n]$ is *ε-regular for A* if for every $I' \subset I, J' \subset J$ satisfying

$$|I'| > \varepsilon|I|, \quad |J'| > \varepsilon|J|$$

we have

$$|\rho_A(I', J') - \rho_A(I, J)| < \varepsilon.$$

The following is a version of the regularity lemma for directed graphs which follows quickly from a stronger result of Alon and Shapira [AS04, Lemma 3.1]. Note that [AS04, Lemma 3.1] is stated for directed graphs without loops, which in the present setting means that it only applies to matrices A with diagonal entries equal to zero. However, Lemma 5.2 follows from applying [AS04, Lemma 3.1] to the matrix A' formed be setting the diagonal entries of A to zero, and noting that the diagonal has a negligible impact on the edge densities $\rho_A(I, J)$ when $|I|, |J| \gg n$.

Lemma 5.2 (Regularity Lemma)**.** *Let $\varepsilon > 0$. There exists $m_0 \in \mathbb{N}$ with $\varepsilon^{-1} \leq m_0 \ll_\varepsilon 1$ such that for all n sufficiently large depending on ε, for every $n \times n$ non-negative matrix A there is a partition of $[n]$ into $m_0 + 1$ sets $I_0, I_1, \ldots, I_{m_0}$ with the following properties:*

(1) $|I_0| < \varepsilon n$;
(2) $|I_1| = |I_2| = \cdots = |I_{m_0}|$;
(3) all but at most εm_0^2 of the pairs (I_k, I_l) are ε-regular for A.

Remark 5.3. The dependence on ε of the bound $m_0 \leq O_\varepsilon(1)$ is very bad: a tower of exponentials of height $O(\varepsilon^{-C})$. Indeed, as in Szemerédi's proof for the setting of bipartite graphs [Sze78], the proof in [AS04] gives such a bound with $C = 5$. It was shown by Gowers that for undirected graphs one cannot do better than $C = 1/16$ in general [Gow97]. As remarked in [AS04], his argument carries over to give a similar result for directed graphs.

We will apply this in Section 5.3 to partition the standard deviation profile into a bounded number of manageable submatrices. The following elementary fact from linear algebra will be used to lift the invertibility properties obtained for these submatrices back to the whole matrix.

Lemma 5.4 (Schur complement bound)**.** *Let $M \in \mathcal{M}_{N+n}(\mathbb{C})$, which we write in block form as*

$$M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$$

for $A \in \mathcal{M}_N(\mathbb{C}), B \in \mathcal{M}_{N,n}(\mathbb{C}), C \in \mathcal{M}_{n,N}(\mathbb{C}), D \in \mathcal{M}_n(\mathbb{C})$. Assume that D is invertible. Then

$$s_{N+n}(M) \geq \left(1 + \frac{\|B\|}{s_n(D)}\right)^{-1} \left(1 + \frac{\|C\|}{s_n(D)}\right)^{-1} \min\Big(s_n(D),\ s_N(A - BD^{-1}C)\Big). \tag{5.1}$$

Proof. From the identity

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{pmatrix} I_N & BD^{-1} \\ 0 & I_n \end{pmatrix} \begin{pmatrix} A - BD^{-1}C & 0 \\ 0 & D \end{pmatrix} \begin{pmatrix} I_N & 0 \\ D^{-1}C & I_n \end{pmatrix}$$

we have

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix}^{-1} = \begin{pmatrix} I_N & 0 \\ -D^{-1}C & I_n \end{pmatrix} \begin{pmatrix} (A - BD^{-1}C)^{-1} & 0 \\ 0 & D^{-1} \end{pmatrix} \begin{pmatrix} I_N & -BD^{-1} \\ 0 & I_n \end{pmatrix}.$$

We can use the triangle inequality to bound the operator norm of the first and third matrices on the right hand side by $1+\|BD^{-1}\|$ and $1+\|CD^{-1}\|$, respectively. Now by sub-multiplicativity of the operator norm,

$$\begin{aligned} \|M^{-1}\| &\leq (1 + \|BD^{-1}\|)(1 + \|D^{-1}C\|) \max(\|(A - BD^{-1}C)^{-1}\|, \|D^{-1}\|) \\ &\leq \left(1 + \frac{\|B\|}{s_n(D)}\right)\left(1 + \frac{\|C\|}{s_n(D)}\right) \max(\|(A - BD^{-1}C)^{-1}\|, \|D^{-1}\|). \end{aligned}$$

The bound (5.1) follows after taking reciprocals. □

5.2. **Control on the operator norm.** The following lemma summarizes the control we will need on the operator norm of submatrices and products of submatrices of M.

Lemma 5.5 (Control on the operator norm)**.** *Let $\xi \in \mathbb{C}$ be a centered random variable with $\mathbb{E}\,|\xi|^{4+\eta} \leq 1$ for some $\eta \in (0,1)$. Let $\theta \in (0,1)$. Then the following hold for all $n \geq 1$:*

(a) (Control for sparse matrices) If $A \in \mathcal{M}_n([0,1])$ is a fixed matrix and $X = (\xi_{ij})$ is an $n \times n$ matrix of iid copies of ξ, then

$$\|A \circ X\| \ll \tau\sqrt{n} \tag{5.2}$$

except with probability $O_\tau(n^{-\eta/8})$, where $\tau = \tau(A) \in [0,1]$ is any number such that

$$\sum_{k=1}^n a_{ik}^2,\ \sum_{k=1}^n a_{kj}^2 \leq \tau^2 n \tag{5.3}$$

for all $i, j \in [n]$, and

$$\sum_{i,j=1}^n a_{ij}^4 \leq \tau^4 n^2. \tag{5.4}$$

(b) (Control for matrix products) Let $m \in [\theta n, n]$. If $A \in \mathcal{M}_{n,m}([0,1])$ and $D \in \mathcal{M}_{m,n}(\mathbb{C})$ are fixed matrices with $\|D\| \leq 1$, and $X = (\xi_{ij})$ is an $n \times m$ matrix of iid copies of ξ, then

$$\|D(A \circ X)\| \ll_\eta \sqrt{m} \tag{5.5}$$

except with probability $O_\theta(n^{-\eta/8})$.

Remark 5.6. The probability bounds in the above lemma can be improved under higher moment assumptions on ξ, and improve to exponential bounds under the assumption that ξ is sub-Gaussian (see (1.2)).

We will use standard truncation arguments to deduce Lemma 5.5 from the following bounds on the expected operator norm of random matrices due to Latała and Vershynin.

Theorem 5.7 (Latała [Lat05])**.** *Let n, m be sufficiently large and let Y be an $n \times m$ random matrix with independent, centered entries $Y_{ij} \in \mathbb{R}$ having finite fourth moment. Then*

$$\mathbb{E}\,\|Y\| \ll \max_{i\in[n]} \left(\sum_{j=1}^{m} \mathbb{E}\,Y_{ij}^2\right)^{1/2} + \max_{j\in[m]} \left(\sum_{i=1}^{n} \mathbb{E}\,Y_{ij}^2\right)^{1/2} + \left(\sum_{i=1}^{n}\sum_{j=1}^{m} \mathbb{E}\,Y_{ij}^4\right)^{1/4}. \tag{5.6}$$

Theorem 5.8 (Vershynin [Ver11])**.** *Let $\eta \in (0,1)$ and n, m, N sufficiently large natural numbers. Let $D \in \mathcal{M}_{m,N}(\mathbb{R})$ be a deterministic matrix satisfying $\|D\| \leq 1$ and $Y \in \mathcal{M}_{N,n}(\mathbb{R})$ be a random matrix with independent centered entries Y_{ij} satisfying $\mathbb{E}\,|Y_{ij}|^{4+\eta} \leq 1$. Then*

$$\mathbb{E}\,\|DY\| \ll_\eta \sqrt{n} + \sqrt{m}. \tag{5.7}$$

Proof of Lemma 5.5. We begin with (a). By splitting X into real and imaginary parts and applying the triangle inequality we may assume ξ is a real-valued random variable. Set $\eta_0 = \min(1/4, \eta/32)$ and define the product event

$$\mathcal{E} = \bigwedge_{i,j=1}^{n} \mathcal{E}_{ij}; \qquad \mathcal{E}_{ij} = \big\{|\xi_{ij}| \leq n^{1/2-\eta_0}\big\}. \tag{5.8}$$

By Markov's inequality,

$$\mathbb{P}(\mathcal{E}_{ij}^c) \leq n^{-(4+\eta)(1/2-\eta_0)} \leq n^{-1} \tag{5.9}$$

for all $i, j \in [n]$. By the union bound,

$$\mathbb{P}(\mathcal{E}^c) \leq n^2 n^{-(4+\eta)(1/2-\eta_0)} \leq n^{-\eta/8}. \tag{5.10}$$

We denote

$$X' = (\xi'_{ij}) = (\xi_{ij} - \mathbb{E}\,\xi_{ij}\,\mathbb{1}_{\mathcal{E}_{ij}}) = X - \mathbb{E}(X\,\mathbb{1}_{\mathcal{E}}).$$

First we show

$$\|A \circ \mathbb{E}(X\,\mathbb{1}_{\mathcal{E}})\| \ll \tau\sqrt{n}. \tag{5.11}$$

Since the variables ξ_{ij} are centered,

$$|\,\mathbb{E}(\xi_{ij}\,\mathbb{1}_{\mathcal{E}_{ij}})| = |\,\mathbb{E}(\xi_{ij}\,\mathbb{1}_{\mathcal{E}_{ij}^c})|. \tag{5.12}$$

By two applications of Hölder's inequality and (5.9),

$$|\,\mathbb{E}(\xi_{ij}\,\mathbb{1}_{\mathcal{E}_{ij}^c})| \le (\mathbb{E}\,|\xi_{ij}|^4)^{1/4}\,\mathbb{P}(\mathcal{E}_{ij}^c)^{3/4} \le n^{-3/4}.$$

Thus,

$$\|A \circ \mathbb{E}(X\,\mathbb{1}_{\mathcal{E}})\| \le \|A \circ \mathbb{E}(X\,\mathbb{1}_{\mathcal{E}})\|_{\mathrm{HS}} \le n^{-3/4}\|A\|_{\mathrm{HS}} \le \tau n^{1/4} \tag{5.13}$$

which yields (5.11) with room to spare.

Now from (5.10), (5.11) and the triangle inequality it is enough to show

$$\mathbb{P}\big(\mathcal{E} \wedge \big\{\|A \circ X'\| \ge C\tau\sqrt{n}\big\}\big) = O_\tau(n^{-\eta/8}) \tag{5.14}$$

for a sufficiently large constant $C > 0$ (we will actually show an exponential bound). First note that the variables $\xi'_{ij}\,\mathbb{1}_{\mathcal{E}_{ij}}$ are centered and satisfy $\mathbb{E}\,|\xi'_{ij}\,\mathbb{1}_{\mathcal{E}_{ij}}|^4 = O(1)$. It follows from Theorem 5.7 that

$$\begin{aligned}\mathbb{E}\,\mathbb{1}_{\mathcal{E}}\,\|A \circ X'\| &\ll \max_{i\in[n]}\left(\sum_{j=1}^n a_{ij}^2\right)^{1/2} + \max_{j\in[n]}\left(\sum_{i=1}^n a_{ij}^2\right)^{1/2} + \left(\sum_{i,j=1}^n a_{ij}^4\right)^{1/4} \\ &\ll \tau\sqrt{n}.\end{aligned}$$

Thus, (5.14) will follow if we can show

$$\mathbb{P}\big(\|A \circ X'\|\,\mathbb{1}_{\mathcal{E}} - \mathbb{E}\,\|A \circ X'\|\,\mathbb{1}_{\mathcal{E}} \ge \tau\sqrt{n}\big) = O_\tau(n^{-\eta/8}). \tag{5.15}$$

This in turn follows in a routine manner from Talagrand's inequality [Tal96, Theorem 6.6] (see also [AGZ10, Corollary 4.4.11]): Observe that $X \mapsto \|A \circ X\|$ is a convex and 1-Lipschitz function on the space $\mathcal{M}_n(\mathbb{R})$ equipped with the (Euclidean) Hilbert–Schmidt metric. Since the matrix $X'\,\mathbb{1}_{\mathcal{E}}$ has centered entries that are bounded by $O(n^{1/2-\eta_0})$, Talagrand's inequality gives that the left hand side of (5.15) is bounded by

$$O\big(\exp(-c\tau^2 n/(n^{1/2-\eta_0})^2)\big) = O\big(\exp(-c\tau^2 n^{2\eta_0})\big) \tag{5.16}$$

which gives (5.15) with plenty of room.

Now we turn to part (b). The proof follows a very similar truncation argument to the one in part (a), so we only indicate the necessary modifications. As before, by splitting D and X into real and imaginary parts and applying the triangle inequality we may assume D and X are real matrices. We define $\mathcal{E}$ as in (5.8), with

$$\mathcal{E}_{ij} = \big\{|\xi_{ij}| \le (n\sqrt{m})^{1/3-\eta_1}\big\} \tag{5.17}$$

and

$$\eta_1 = \frac{1}{4}\frac{\eta}{4+\eta}. \tag{5.18}$$

With this choice of η_1, Markov's inequality and the union bound give $\mathbb{P}(\mathcal{E}^c) = O_\theta(n^{-\eta/8})$. Taking $X' = X - \mathbb{E}(X\,\mathbb{1}_{\mathcal{E}})$ as before, we can bound

$$\|D(A \circ \mathbb{E}(X\,\mathbb{1}_{\mathcal{E}}))\| \le \|A \circ \mathbb{E}(X\,\mathbb{1}_{\mathcal{E}})\|$$

by submultiplicativity of the operator norm, and the same argument as before gives

$$\|A \circ \mathbb{E}(X\,\mathbb{1}_{\mathcal{E}})\| \le nm(n\sqrt{m})^{-\frac{3}{4}(4+\eta)(1/3-\eta_1)} = m^{1/2-\eta/32} = o(\sqrt{m}). \tag{5.19}$$

Since $X'\,\mathbb{1}_{\mathcal{E}}$ has centered entries with finite moments of order $4+\eta$, by Theorem 5.8 we have

$$\mathbb{E}\,\|D(A \circ X'\,\mathbb{1}_{\mathcal{E}})\| \ll_\eta \sqrt{m}. \tag{5.20}$$

The mapping $X \mapsto \|D(A \circ X)\|$ is convex and 1-Lipschitz with respect to the Hilbert–Schmidt metric on $\mathcal{M}_n(\mathbb{R})$ (since $\|D\| \le 1$) so using Talagrand's inequality as in part (a) we find that

$$\begin{aligned}\mathbb{P}\big(\|D(A \circ X'\,\mathbb{1}_{\mathcal{E}})\| - \mathbb{E}\,\|D(A \circ X'\,\mathbb{1}_{\mathcal{E}})\| \ge \sqrt{m}\big) &\ll \exp\Big(-cm/(n\sqrt{m})^{2/3-2\eta_1}\Big)\\ &\le \exp\big(-c'(\theta)n^{c\eta}\big)\end{aligned}$$

for some constant $c > 0$ and $c'(\theta) > 0$ sufficiently small depending on θ. As the last line is bounded by $O_\theta(n^{-\eta/8})$, the result follows from the above, (5.19), (5.20) and the triangle inequality by the same argument as for part (a). □

5.3. **Decomposition of the standard deviation profile.** We now begin the proof of Theorem 1.19, which occupies the remainder of the paper. In the present subsection we prove Lemma 5.9 below, which shows that the standard deviation profile A can be partitioned into a bounded collection of submatrices with certain nice properties. For the motivation behind this lemma (and the notation $J_{\text{free}}, J_{\text{cyc}}$) see Section 1.5.

Lemma 5.9. *Let A be an $n \times n$ matrix with entries $a_{ij} \in [0,1]$. Let $\varepsilon, \delta, \sigma_0 \in (0,1)$, and assume ε is sufficiently small depending on δ. There exists $0 \le m \ll_\varepsilon 1$, a partition*

$$\begin{aligned}[n] &= J_{\text{bad}} \cup J_{\text{free}} \cup J_{\text{cyc}}\\ &= J_{\text{bad}} \cup J_{\text{free}} \cup J_1 \cup \cdots \cup J_m\end{aligned} \tag{5.21}$$

and a set $F \subset [n]^2$ satisfying the following properties:

(1) $\varepsilon n \ll |J_{\text{bad}}| \ll \delta^{1/2} n$.
(2) $|F| \ll \delta n^2$, and for all $i \in J_{\text{free}}$,

$$|\{j \in J_{\text{free}} : (i,j) \in F\}|,\ |\{j \in J_{\text{free}} : (j,i) \in F\}| \le \delta^{1/2} n. \tag{5.22}$$

(3) If $J_{\text{free}} \ne \varnothing$ then there is a permutation $\tau : J_{\text{free}} \to J_{\text{free}}$ such that for all $(i,j) \in J_{\text{free}} \times J_{\text{free}} \setminus F$ with $\tau(i) \ge \tau(j)$, $a_{ij} < \sigma_0$.

(4) If $m \geq 1$ then

$$|J_1| = \cdots = |J_m| \gg_\varepsilon n \tag{5.23}$$

and there is a permutation $\pi : [m] \to [m]$ such that for all $1 \leq k \leq m$, $A(\sigma_0)_{J_k, J_{\pi(k)}}$ is $(2\delta, 2\varepsilon)$-super-regular.

Proof. We begin by applying Lemma 5.2 to $A(\sigma_0)$ to obtain $m_0 \in \mathbb{N}$ with $\varepsilon^{-1} \leq m_0 = O_\varepsilon(1)$ and a partition $[n] = I_0 \cup \cdots \cup I_{m_0}$ satisfying the properties in that lemma.

The partition $I_0, \dots, I_{m_0}$ is almost what we need. In the remainder of the proof we perform a "cleaning" procedure (as it is commonly referred to in the extremal combinatorics literature) to obtain a partition $J_0, \dots, J_{m_0}$ with improved properties, where $J_k \subset I_k$ for each $1 \leq k \leq m_0$, and $J_0 \supset I_0$ collects the leftover elements.

We start by forming a *reduced digraph* $\mathcal{R} = ([m_0], E)$ on the vertex set $[m_0]$ with directed edge set

$$E := \Big\{(k,l) \in [m_0]^2 : (I_k, I_l) \text{ is } \varepsilon\text{-regular and } \rho_{A(\sigma_0)}(I_k, I_l) > 5\delta\Big\}. \tag{5.24}$$

Next we find a (possibly empty) set $T \subset [m_0]$ such that the induced subgraph $\mathcal{R}(T)$ is covered by vertex-disjoint directed cycles, and the induced subgraph $\mathcal{R}([m_0] \setminus T)$ is cycle-free. Such a set can be obtained by greedily removing cycles and the associated vertices from $\mathcal{R}$ until the remaining graph has no more directed cycles. By relabeling $I_1, \dots, I_{m_0}$ we may take $T = [m]$, where $m \in [0, m_0]$.

Assuming $m \neq 0$, the fact that $\mathcal{R}([m])$ is cycle-free is equivalent to the existence of a permutation $\pi : [m] \to [m]$ such that $(k, \pi(k)) \in E$ for all $1 \leq k \leq m$. Now we will obtain the sets $J_1, \dots, J_m$ obeying the properties in part (4) of the lemma. Let $1 \leq k \leq m$. Since $(I_k, I_{\pi(k)})$ is ε-regular with density $\rho_k := \rho_{A(\sigma_0)}(I_k, I_{\pi(k)}) > 5\delta$, if we assume $\varepsilon \leq \delta$ then for every $I \subset I_k, J \subset I_{\pi(k)}$ with $|I|, |J| \geq \varepsilon |I_k|$,

$$e_{A(\sigma_0)}(I, J) \geq (\rho_k - \varepsilon)|I||J| \geq 4\delta |I||J|. \tag{5.25}$$

Letting

$$I_k' = \big\{i \in I_k : |\mathcal{N}_{A(\sigma_0)}(i) \cap I_{\pi(k)}| < 4\delta |I_k|\big\}$$

we have $e_{A(\sigma_0)}(I_k', I_{\pi(k)}) < 4\delta |I_k'||I_{\pi(k)}|$, and it follows that $|I_k'| \leq \varepsilon |I_k|$. Similarly, letting

$$I_k'' = \big\{i \in I_k : |\mathcal{N}_{A(\sigma_0)^\intercal}(i) \cap I_{\pi^{-1}(k)}| < 4\delta |I_k|\big\}$$

we have $|I_k''| \leq \varepsilon |I_k|$. Letting $I_k^* \subset I_k$ be a set of size $\lfloor 2\varepsilon |I_k| \rfloor$ containing $I_k' \cup I_k''$, we take

$$J_k = I_k \setminus I_k^*. \tag{5.26}$$

With this definition we have $|J_1| = \cdots |J_m|$, and for each $1 \leq k \leq m, i \in J_k$,

$$|\mathcal{N}_{A(\sigma_0)}(i) \cup J_{\pi(k)}|,\ |\mathcal{N}_{A(\sigma_0)^\intercal}(i) \cap J_{\pi^{-1}(k)}| \geq (4\delta - 2\varepsilon)|I_k| \geq 2\delta |J_k|. \tag{5.27}$$

Furthermore, for each $1 \leq k \leq m$ and $I \subset J_k, J \subset J_{\pi(k)}$ with $|I|, |J| \geq 2\varepsilon |J_k|$, if we assume $\varepsilon \leq 1/4$ then $|I|, |J| \geq \varepsilon |I_k|$, so by (5.25)

$$e_{A(\sigma_0)}(I, J) \geq 4\delta |I||J|. \tag{5.28}$$

It follows that for every $1 \le k \le m$ the submatrix $A(\sigma_0)_{J_k,J_{\pi(k)}}$ is $(2\delta, 2\varepsilon)$-super-regular, which concludes the proof of part (4) of the lemma.

Now we prove parts (2) and (3). We will obtain J_{free} by removing a small number of bad elements from $I_{\text{free}} := \bigcup_{k=m+1}^{m_0} I_k$. Since the induced subgraph $\mathcal{R}([m+1, m_0])$ is cycle-free we may relabel $I_{m+1}, \dots, I_{m_0}$ so that

$$(k,l) \notin E \text{ for all } m < l \le k \le m_0. \tag{5.29}$$

We take

$$F = \big\{(i,j) \in [n]^2 : (i,j) \in I_k \times I_l \text{ for some } (k,l) \notin E\big\}. \tag{5.30}$$

The contribution to F from irregular pairs (I_k, I_l) is at most εn^2 by the regularity of the partition $I_0, \dots, I_{m_0}$, and the contribution from pairs (I_k, I_l) with density less than 5δ is at most $5\delta n^2$. Hence,

$$|F| \le \varepsilon n^2 + 5\delta n^2 \le 6\delta n^2 \tag{5.31}$$

giving the first estimate in (2) (recall that we assumed $\varepsilon \le \delta$). Setting

$$I'_{\text{free}} = \Big\{ i \in I_{\text{free}} : \max\big(|\{j \in [n] : (i,j) \in F\}|, |\{j \in [n] : (j,i) \in F\}|\big) \ge \delta^{1/2} n \Big\} \tag{5.32}$$

it follows from (5.31) that

$$|I'_{\text{free}}| \le 12\delta^{1/2} n. \tag{5.33}$$

Let $I^*_{\text{free}} \subset I_{\text{free}}$ be any set containing I'_{free} of size $\min(|I_{\text{free}}|, \lfloor 12\delta^{1/2} n \rfloor)$ and take $J_{\text{free}} = I_{\text{free}} \setminus I^*_{\text{free}}$. The bounds (5.22) now follow immediately from (5.32). For part (3), from (5.29) we may take for τ any ordering of the elements of J_{free} that respects the order of the sets $J_k := I_k \setminus I^*_{\text{free}}$, i.e. so that $\tau(j) \ge \tau(i)$ for all $i \in J_k, j \in J_l$ and all $m < l \le k \le m_0$.

Finally, taking

$$J_{\text{bad}} = I_0 \cup I^*_{\text{free}} \cup \bigcup_{k=1}^{m} I^*_k. \tag{5.34}$$

we have

$$|J_{\text{bad}}| \le \varepsilon n + 12\delta^{1/2} n + 2\varepsilon n \le 15\delta^{1/2} n$$

giving the upper bound in part (1). Now recalling that we took $|I^*_{\text{free}}| = 12\delta^{1/2}$ and $|I^*_k| = \lfloor 2\varepsilon |I_k| \rfloor$ for all $1 \le k \le m$, we also have the lower bound

$$\begin{aligned}
|J_{\text{bad}}| &\ge \min\left(|I^*_{\text{free}}|, \Big| \bigcup_{k=1}^{m} I^*_k \Big| \right) \\
&\ge \min\left(\lfloor 12\delta^{1/2} n \rfloor, |I_{\text{free}}|, 2\varepsilon \Big| \bigcup_{k=1}^{m} I_k \Big| - m \right) \\
&= \min\left(\lfloor 12\delta^{1/2} n \rfloor, \Big| \bigcup_{k=m+1}^{m_0} I_k \Big|, 2\varepsilon \Big| \bigcup_{k=1}^{m} I_k \Big| - m \right) \\
&\gg \varepsilon n
\end{aligned}$$

where we used that at least one of the sets $I_{\mathrm{free}} = \bigcup_{k=m+1}^{m_0} I_k$, $I_{\mathrm{cyc}} = \bigcup_{k=1}^{m} I_k$ must be of size at least $n/4$, say. This gives the lower bound in part (1) and completes the proof. □

5.4. **High level proof of Theorem 1.19.** In this subsection we prove Theorem 1.19 on two lemmas (Lemmas 5.10 and 5.11) which give control on the smallest singular values of the submatrices $M_{J_{\mathrm{free}}}$ and (perturbations of) $M_{J_{\mathrm{cyc}}}$, with $J_{\mathrm{free}}, J_{\mathrm{cyc}}$ as in Lemma 5.9. The proofs of these lemmas are deferred to the remaining subsections.

By our moment assumptions on ξ it follows that ξ is κ_0-spread for some $\kappa_0 = O(\mu_{4+\eta}^2)$ (see Remark 1.3). By Lemma 2.5 and multiplying X and B by a phase we may assume ξ has $O(\mu_{4+\eta}^2)$-controlled second moment. Without loss of generality we may assume $\eta < 1$. We introduce parameters $\sigma_0, \delta, \varepsilon \in (0,1)$, where σ_0 will be chosen sufficiently small depending on $r_0, \mu_{4+\eta}$, δ will be taken small depending on $r_0, \eta, \mu_{4+\eta}$, and ε will be taken sufficiently small depending on σ_0, δ. Below we introduce an additional free parameter $\theta \in (0,1)$ to be chosen later sufficiently small depending on ε. For the remainder of the proof we assume that n is sufficiently large depending on all parameters (which will only depend on r_0, K_0, η and $\mu_{4+\eta}$).

We begin by summarizing the control we have on the operator norm of submatrices of $A \circ X$ due to Lemma 5.5. By Lemma 5.5(a) with $\tau = 1$ we may restrict the sample space to the boundedness event

$$\mathcal{B}(K) = \big\{ \|A \circ X\| \le K\sqrt{n} \big\} \tag{5.35}$$

for some $K = O(\mu_{4+\eta})$. In particular, for all $I, J \subset [n]$ we now have the crude bound

$$\|(A \circ X)_{I,J}\| \le K\sqrt{n}. \tag{5.36}$$

We also note that by Lemma 5.5(a) we have that for any fixed $B = (b_{ij}) \in \mathcal{M}_n([0,1])$,

$$\mathbb{P}\Big(\|(B \circ X)_{I,J}\| \le \tau K \sqrt{|I|} \Big) = 1 - O_\tau(n^{-\eta/8}) \tag{5.37}$$

for all $I, J \subset [n]$ with $|I| = |J|$, and any $\tau \le 1$ satisfying

$$\tau \ge \max\left(\max_{i \in I} \left(\frac{1}{n} \sum_{j \in J} b_{ij}^2 \right)^{1/2}, \max_{j \in J} \left(\frac{1}{n} \sum_{i \in I} b_{ij}^2 \right)^{1/2}, \left(\frac{1}{n} \sum_{i,j=1}^{n} b_{ij}^4 \right)^{1/4} \right). \tag{5.38}$$

Note that for $B = A$, (5.37) gives improved control on $\|(A \circ X)_{I,J}\|$ for submatrices that are either small or sparse, but only holding with high probability for fixed I, J. (We state (5.37) for general $B \in \mathcal{M}_n([0,1])$ as at one point we will apply this to a residual matrix obtained by subtracting off a collection of "bad" entries from A.)

Finally, let $\theta \in (0,1)$ to be chosen later depending on ε. From Lemma 5.5(b) we have that for every $I, J \subset [n]$ with $|I| \geq |J| \geq \theta n$ and every deterministic $|J| \times |I|$ matrix D,

$$\mathbb{P}\Big(\|D(A \circ X)_{I,J}\| \leq K'\|D\|\sqrt{|J|}\Big) = 1 - O_\theta(|I|^{-\eta/8}) \tag{5.39}$$

for some K' sufficiently large depending on η and $\mu_{4+\eta}$.

We now apply Lemma 5.9 (assuming ε is sufficiently small depending on δ) to obtain a partition $[n] = J_{\mathrm{bad}} \cup J_{\mathrm{free}} \cup J_{\mathrm{cyc}}$ and a set $F \subset [n]^2$ satisfying the properties (1)–(4) in the lemma. In the following we abbreviate $M_{\mathrm{free}} := M_{J_{\mathrm{free}}}$ and $M_{\mathrm{cyc}} := M_{J_{\mathrm{cyc}}}$.

Lemma 5.10. *Assume $n_1 := |J_{\mathrm{free}}| \geq \delta^{1/2} n$. If σ_0, δ are sufficiently small depending on r_0 and K (from (5.35)), then*

$$s_{n_1}(M_{\mathrm{free}}) \gg_{K,r_0} \sqrt{n} \tag{5.40}$$

except with probability $O_{K,r_0,\delta}(n^{-\eta/9})$.

Lemma 5.11. *Assume $n_2 := |J_{\mathrm{cyc}}| \geq \delta^{1/2} n$. Fix $\gamma > 1/2$ and let $W \in \mathcal{M}_{n_2}(\mathbb{C})$ be a deterministic matrix with $\|W\| \leq n^\gamma$. There exists $\beta = \beta(\gamma, \sigma_0, \delta)$ such that if $\varepsilon = \varepsilon(\sigma_0, \delta)$ is sufficiently small,*

$$\mathbb{P}\Big(s_{n_2}(M_{\mathrm{cyc}} + W) \leq n^{-\beta}\Big) \ll_{K_0,K,\gamma,\delta,\sigma_0,\mu_{4+\eta}} \sqrt{\frac{\log n}{n}}. \tag{5.41}$$

Remark 5.12. We note that in the proof of Lemma 5.11 we do not make use of the fact that the atom variable ξ has more than two finite moments (the dependence on $\mu_{4+\eta}$ is only through the parameter $\kappa_0 = O(\mu_{4+\eta}^2)$). In particular, we can remove the extra moment hypotheses in Theorem 1.19 under the additional assumption that the standard deviation profile A contains a generalized diagonal of block submatrices which are super-regular and of dimension linear in n (that is, if we can take $J_{\mathrm{bad}} = J_{\mathrm{free}} = \varnothing$ in (5.21)).

We defer the proofs of Lemmas 5.10 and 5.11 to subsequent sections, and conclude the proof of Theorem 1.19. We proceed in the following steps:

Step 1: Bound the smallest singular value of M_{free} using Lemma 5.10.

Step 2: Bound the smallest singular value of

$$M_1 := M_{J_{\mathrm{free}} \cup J_{\mathrm{bad}},\, J_{\mathrm{free}} \cup J_{\mathrm{bad}}} = \begin{pmatrix} M_{\mathrm{free}} & B_1 \\ C_1 & M_0 \end{pmatrix}. \tag{5.42}$$

using the result of Step 1, the Schur complement bound of Lemma 5.4, and the bounds (5.37), (5.39).

Step 3: Bound the smallest singular value of

$$M = \begin{pmatrix} M_{\mathrm{cyc}} & B_2 \\ C_2 & M_1 \end{pmatrix}. \tag{5.43}$$

using the result of Step 2, the Schur complement bound of Lemma 5.4, and Lemma 5.11.

The case that one of J_{free} or J_{cyc} is small (or empty) can be handled essentially by skipping either Step 1 or Step 3. We will begin by assuming

$$|J_{\mathrm{free}}|,\, |J_{\mathrm{cyc}}| \geq \delta^{1/2} n \tag{5.44}$$

and address the case that this does not hold at the end.

Step 1. By Lemma 5.10 and the assumption (5.44), we can take σ_0 and δ sufficiently small depending on r_0 and K (where K depends only $\mu_{4+\eta}$) such that

$$s_{\min}(M_{\mathrm{free}}) \gg_{K,r_0} \sqrt{n} \tag{5.45}$$

except with probability $O_{K,r_0,\delta}(n^{-\eta/9})$. We now fix $\sigma_0 = \sigma_0(r_0, \mu_{4+\eta})$ once and for all, but leave δ free to be taken smaller if necessary. By independence of the entries of M we may now condition on a realization of M_{free} such that (5.45) holds.

Step 2. From Lemma 5.4, (5.36) and (5.45),

$$\begin{aligned} s_{\min}(M_1) &\gg \left(1 + \frac{K\sqrt{n}}{s_{\min}(M_{\mathrm{free}})}\right)^{-2} \min\left[s_{\min}(M_{\mathrm{free}}), s_{\min}(M_0 - C_1 M_{\mathrm{free}}^{-1} B_1)\right] \\ &\gg_{K,r_0} \min\left[\sqrt{n}, s_{\min}(M_0 - C_1 M_{\mathrm{free}}^{-1} B_1)\right]. \end{aligned} \tag{5.46}$$

Again by (5.36) and (5.45),

$$\|C_1 M_{\mathrm{free}}^{-1}\| \leq \frac{\|C_1\|}{s_{\min}(M_{\mathrm{free}})} \ll_{K,r_0} 1. \tag{5.47}$$

Since C_1 is independent of B_1, we may condition on it (recall that we already conditioned on M_{free}) and apply (5.39) to conclude

$$\|C_1 M_{\mathrm{free}}^{-1} B_1\| \leq K' \|C_1 M_{\mathrm{free}}^{-1}\| |J_{\mathrm{bad}}|^{1/2} \ll_{K,K',r_0} |J_{\mathrm{bad}}|^{1/2} \tag{5.48}$$

except with probability $O_\theta(n_1^{-\eta/8}) = O_{\delta,\varepsilon}(n^{-\eta/9})$, where we have used the lower bound $|J_{\mathrm{bad}}| \gg \varepsilon n$ from Lemma 5.9(1) and fixed θ sufficiently small depending on ε. On the other hand, by from the triangle inequality and the bound (5.37) with $\tau = 1$,

$$s_{\min}(M_0) = s_{\min}(Z_{J_{\mathrm{bad}}}\sqrt{n} + (A \circ X)_{J_{\mathrm{bad}}}) \geq r_0\sqrt{n} - K|J_{\mathrm{bad}}|^{1/2} \tag{5.49}$$

except with probability $O(|J_{\mathrm{bad}}|^{-\eta/8}) = O_\varepsilon(n^{-\eta/9})$. Combining the previous two displays,

$$s_{\min}(M_0 - C_1 M_{\mathrm{free}}^{-1} B_1) \geq r_0\sqrt{n} - O_{K,K',r_0}(|J_{\mathrm{bad}}|^{1/2}) \tag{5.50}$$

except with probability $O_{\delta,\varepsilon}(n^{-\eta/9})$. Since $|J_{\mathrm{bad}}| \ll \delta^{1/2}n$ we can take δ smaller, if necessary, depending on r_0, K, K' to conclude that

$$s_{\min}(M_0 - C_1 M_{\mathrm{free}}^{-1} B_1) \geq (r_0/2)\sqrt{n} \tag{5.51}$$

except with probability $O_{\delta,\varepsilon}(n^{-\eta/9})$. We may henceforth condition on the event that (5.51) holds. Together with (5.46) this implies

$$s_{\min}(M_1) \gg_{K,r_0} \sqrt{n}. \tag{5.52}$$

At this point we fix $\delta = \delta(r_0, \eta, \mu_{4+\eta})$.

Step 3. Condition on a realization of M_1 such that (5.52) holds, and also on the matrices B_2, C_2 in (5.43). By (5.36) we have $\|B_2\|, \|C_2\| \leq K\sqrt{n}$. Applying Lemma 5.4,

$$\begin{aligned} s_n(M) &\gg \left(1 + \frac{K\sqrt{n}}{s_{\min}(M_1)}\right)^{-2} \min\left[s_{\min}(M_1),\, s_{\min}(M_{\mathrm{cyc}} - B_2 M_1^{-1} C_2)\right] \\ &\gg_{K,r_0} \min\left[\sqrt{n}, s_{\min}(M_{\mathrm{cyc}} - B_2 M_1^{-1} C_2)\right]. \end{aligned} \tag{5.53}$$

Again by (5.36) and (5.52),

$$\|B_2 M_1^{-1} C_2\| \leq \frac{K^2 n}{s_{\min}(M_1)} \ll_{K,r_0} \sqrt{n} \tag{5.54}$$

(unlike in Step 2, here we did not need the stronger control on matrix products provided by (5.5) from the matrix boundedness property). Now since M_2 is independent of M_1, B_2, C_2, we can apply Lemma 5.11 with $\gamma = 0.51$ (say), fixing ε sufficiently small depending on $\sigma_0(r_0, \mu_{4+\eta})$ and $\delta(r_0, \eta, \mu_{4+\eta})$, to obtain

$$\mathbb{P}\Big(s_{\min}(M_{\mathrm{cyc}} - B_2 M_1^{-1} C_2) \leq n^{-\beta}\Big) \ll_{K_0,r_0,\eta,\mu_{4+\eta}} \sqrt{\frac{\log n}{n}} \tag{5.55}$$

for some $\beta = \beta(r_0, \eta, \mu_{4+\eta}) > 0$, where we have substituted $K = O(\mu_{4+\eta})$. The result now follows from the above and (5.53), taking $\alpha = \min(\eta/9, 1/4)$, say.

It only remains to address the case that the assumption (5.44) fails. We may assume that δ is small enough that only one of these bounds fails. In this case we simply redefine J_{bad} to include the smaller of $J_{\mathrm{cyc}}, J_{\mathrm{free}}$. Note that we still have $|J_{\mathrm{bad}}| = O(\delta^{1/2}n)$. If $|J_{\mathrm{cyc}}| < \delta^{1/2}n$, then with this new definition of J_{bad} we have $M = M_1$, and the desired bound on $s_n(M)$ follows from (5.52) (with plenty of room). If $|J_{\mathrm{free}}| < \delta^{1/2}n$ then we skip Step 2, proceeding with Step 3 using M_0 in place of M_1. The bound (5.52) in this case follows from (5.49) and the bound $|J_{\mathrm{bad}}| \ll \delta^{1/2}n$, taking δ sufficiently small depending on K, r_0. This concludes the proof of Theorem 1.19.

5.5. **Proof of Lemma 5.10.** We denote

$$A_{\mathrm{bad}} = (a_{ij} 1_{(i,j)\in F}). \tag{5.56}$$

By the estimates on F in Lemma 5.9 we can apply (5.37) with $\tau = O(\delta^{1/4})$ to obtain

$$\|(A_{\mathrm{bad}}(\sigma_0) \circ X)_{J_{\mathrm{free}}}\| \ll K\delta^{1/4}\sqrt{n} \tag{5.57}$$

except with probability at most $O_\delta(n_1^{-\eta/8}) = O_\delta(n^{-\eta/9})$. By another application of (5.37) with $\tau = 1$,

$$\big\|\big((A - A(\sigma_0)) \circ X\big)_{J_{\mathrm{free}}}\big\| \ll K\sigma_0\sqrt{n} \tag{5.58}$$

except with probability at most $O_\delta(n^{-\eta/9})$. Let

$$\widetilde{M}_{\mathrm{free}} := (\widetilde{A} \circ X)_{J_{\mathrm{free}}} + Z_{J_{\mathrm{free}}}\sqrt{n}, \qquad \widetilde{A} := A(\sigma_0) - A_{\mathrm{bad}}(\sigma_0). \tag{5.59}$$

By the above estimates and the triangle inequality,

$$s_{\min}(M_{\mathrm{free}}) \geq s_{\min}(\widetilde{M}_{\mathrm{free}}) - \|((A - \widetilde{A}) \circ X)_{J_{\mathrm{free}}}\| \tag{5.60}$$

$$\geq s_{\min}(\widetilde{M}_{\mathrm{free}}) - O(\delta^{1/4} + \sigma_0)K\sqrt{n} \tag{5.61}$$

except with probability $O_\delta(n^{-\eta/9})$. Thus, it suffices to show

$$s_{\min}(\widetilde{M}_{\mathrm{free}}) \gg_{K,r_0} \sqrt{n}. \tag{5.62}$$

except with probability $O_{K,r_0,\delta}(n^{-\eta/9})$ – the result will then follow from (5.62) and (5.61) by taking δ, σ_0 sufficiently small depending on K, r_0. Furthermore, by Lemma 5.9(3) and conjugating M_{free} by a permutation matrix we may assume that $\widetilde{A}$ is (strictly) upper triangular. Now it suffices to prove the following:

Lemma 5.13. *Let $M = A \circ X + B$ be an $n \times n$ matrix as in Definition 1.4, and further assume that for some $r_0 > 0, K \geq 1, \alpha > 0$,*

- *A is upper triangular;*
- *$B = Z\sqrt{n} = \mathrm{diag}(z_i\sqrt{n})_{i=1}^n$ with $|z_i| \geq r_0$ for all $1 \leq i \leq n$;*
- *ξ is such that for all $n' \geq 1$ and any fixed $A' \in \mathcal{M}_{n'}([0,1])$, $\|A' \circ X'\| \leq K\sqrt{n'}$ except with probability $O((n')^{-\alpha})$.*

Then $s_n(M) \gg_{K,r_0} \sqrt{n}$ except with probability $O_{K,r_0}(1)^\alpha n^{-\alpha}$.

Remark 5.14. The proof gives an implied constant of order $\exp(-O(K/r_0)^{O(1)})$ in the lower bound on $s_n(M)$.

To deduce Lemma 5.10 we apply the above lemma with $M = \widetilde{M}_{\mathrm{free}}$, $\alpha = \eta/8$, K as in (5.35) and $n_1 \gg_\delta n$ in place of n, which gives that (5.62) holds with probability

$$1 - O_{K,r_0}(n_1^{-\eta/8}) = 1 - O_{K,r_0,\delta}(n^{-n/9}) \tag{5.63}$$

where in the first bound we applied our assumption that $\eta < 1$.

Proof. First we note that we may take n to be a dyadic integer, i.e. $n = 2^q$ for some $q \in \mathbb{N}$. Indeed, if this is not the case, then letting 2^q be the smallest dyadic integer larger than n we can increase the dimension of M to 2^q by padding A out with rows and columns of zeros, adding additional rows and columns of iid copies of ξ to X, and extending the diagonal of Z with entries $z_i \equiv r_0$ for $n < i \leq 2^q$. The hypotheses on A and Z in the lemma are still satisfied, and the smallest singular value of the new matrix is a lower bound for that of the original matrix (since the original matrix is a submatrix of the new matrix).

Now fix an arbitrary dyadic filtration $\mathcal{F} = \bigcup_{p \geq 0} \{J_s : s \in \{0,1\}^p\}$ of $[n]$, where we view $\{0,1\}^0$ as labeling the trivial partition of $[n]$, consisting only of the empty string $\varnothing$, so that $J_\varnothing = [n]$. Thus, for every $0 \leq p < q$ and every binary string $s \in \{0,1\}^p$, J_s has cardinality $n2^{-p}$ and is evenly partitioned by J_{s0}, J_{s1}. For a binary string s we abbreviate $M_s := M_{J_s}$ and similarly define A_s, X_s, Z_s. We also write $B_s = M_{J_{s0}, J_{s1}}$, so that we have the block decomposition

$$M_s = \begin{pmatrix} M_{s0} & B_s \\ 0 & M_{s1} \end{pmatrix}. \tag{5.64}$$

For $p \geq 1$ define the boundedness event

$$\mathcal{B}^*(p) = \left\{ \|A \circ X\| \leq K\sqrt{n} \right\} \wedge \left\{ \forall s \in \{0,1\}^p,\ \|A_s \circ X_s\| \leq K\sqrt{n2^{-p}} \right\}. \tag{5.65}$$

By our assumption on ξ we have

$$\mathbb{P}(\mathcal{B}^*(p)) \geq 1 - O(n^{-\alpha}) - 2^p O((n2^{-p})^{-\alpha}) = 1 - O(2^{(1+\alpha)p} n^{-\alpha}). \tag{5.66}$$

For arbitrary $s \in \{0,1\}^p$, by the triangle inequality we have that on $\mathcal{B}^*(p)$,

$$s_{\min}(M_s) \geq s_{\min}(Z_s) - \|A_s \circ X_s\| \geq (r_0 - K2^{-p/2})\sqrt{n}.$$

Setting $p_0 = \lfloor 2\log(2K/r_0) \rfloor + 1$ we have that on $\mathcal{B}^*(p_0)$,

$$s_{\min}(M_s) \geq (r_0/2)\sqrt{n} \tag{5.67}$$

for all $s \in \{0,1\}^{p_0}$. For the remainder of the proof we restrict the sample space to the event $\mathcal{B}^*(p_0)$ and will use the Schur complement bound (Lemma 5.4) to show that the desired lower bound on $s_{\min}(M)$ holds deterministically (note that by (5.66) and our choice of p_0, $\mathcal{B}^*(p_0)$ holds with probability $1 - O_{K,r_0}(n^{-\alpha/2})$).

For $0 \leq p \leq p_0$ let

$$\lambda_p = \min_{s \in \{0,1\}^p} \frac{1}{\sqrt{n}} s_{\min}(M_s). \tag{5.68}$$

From (5.67) we have

$$\lambda_{p_0} \geq r_0/2 \tag{5.69}$$

Now let $1 \leq p \leq p_0$ and $s \in \{0,1\}^{p-1}$. By the block decomposition (5.64) and Lemma 5.4,

$$\begin{aligned} s_{\min}(M_s) &\gg \left(1 + \frac{\|B_s\|}{s_{\min}(M_{s0})}\right)^{-1} \min\left(s_{\min}(M_{s0}), s_{\min}(M_{s1})\right) \\ &\geq (1 + K/\lambda_p)^{-1} \lambda_p \sqrt{n} \end{aligned}$$

so $\lambda_{p-1} \gg (1+K/\lambda_p)^{-1}\lambda_p\sqrt{n}$ for all $0 \le p \le p_0$. Applying this iteratively along with (5.69) we conclude $\lambda_0 \gg_{K,r_0} 1$, i.e.

$$s_{\min}(M) \gg_{K,r_0} \sqrt{n} \tag{5.70}$$

as desired. □

5.6. **Proof of Lemma 5.11.** By multiplying M_{cyc} by a permutation matrix we may assume that $A_k := A_{J_k}$ is $(2\delta, 2\varepsilon)$-super-regular for $1 \le k \le m$ (unlike in the proof of Lemma 5.10 the diagonal matrix $Z\sqrt{n}$ plays no special role here). We denote $J_{\le k} = J_1 \cup \cdots \cup J_k$, and for any matrix W of dimension at least $|J_{\le k}|$ we abbreviate

$$W_k = W_{J_k}, \quad W_{\le k} = W_{J_{\le k}}, \quad W_{\le k-1,k} = W_{J_{\le k-1},J_k}, \quad W_{k,\le k-1} = W_{J_k,J_{\le k-1}} \tag{5.71}$$

so that for $2 \le k \le m$ we have the block decomposition

$$W_{\le k} = \begin{pmatrix} W_{\le k-1} & W_{\le k-1,k} \\ W_{k,\le k-1} & W_k \end{pmatrix}. \tag{5.72}$$

Let us denote

$$n' = |J_1| = \cdots = |J_m| \gg_\varepsilon n. \tag{5.73}$$

For $1 \le k \le m-1$, $\beta > 0$ and a fixed $kn' \times kn'$ matrix W, we denote the event

$$\mathcal{E}_k(\beta, W) := \big\{ s_{kn'}(M_{\le k} + W) > n^{-\beta} \big\}. \tag{5.74}$$

Let $\gamma > 1/2$ and fix an arbitrary matrix $W \in \mathcal{M}_{n',n'}(\mathbb{C})$ with $\|W\| \le n^\gamma$. By our restriction to $\mathcal{B}(K)$ and (5.36) we have

$$\|M_1 + W\| \le (K_0 + K)\sqrt{n} + n^\gamma \le 2n^\gamma \tag{5.75}$$

if n is sufficiently large depending on K_0, K and γ. By Theorem 1.25 there exists $\beta_1(\gamma) = O(\gamma^2)$ such that if ε is sufficiently small depending on σ_0, δ, then

$$\mathbb{P}\left(\mathcal{E}_1(\beta_1, W)^c\right) \ll_{K_0,K,\gamma,\delta,\sigma_0,\varepsilon,\mu_{4+\eta}} \sqrt{\frac{\log n}{n}} \tag{5.76}$$

where we have used (5.73) to write n in $n^{-\beta_1}$ rather than n', and the fact that the atom variable is $O(\mu_{4+\eta}^2)$-spread.

Now let $2 \le k \le m$, and suppose we have found a function $\beta_{k-1}(\gamma)$ such that for any $\gamma > 1/2$ and any fixed $(k-1)n' \times (k-1)n'$ matrix W with $\|W\| \le n^\gamma$,

$$\mathbb{P}(\mathcal{E}_{k-1}(\beta_{k-1}(\gamma), W)^c) \ll_{K_0,K,\gamma,\delta,\sigma_0,\varepsilon,\mu_{4+\eta}} \sqrt{\frac{\log n}{n}}. \tag{5.77}$$

Fix a $kn' \times kn'$ matrix W with $\|W\| \le n^\gamma$. By Lemma 5.4 we have

$$\begin{aligned} s_{kn'}(M_{\le k} + W) \gg \left(1 + \frac{\|(M+W)_{\le k-1,k}\|}{s_{(k-1)n'}(M_{\le k-1} + W_{\le k-1})}\right)^{-1} \left(1 + \frac{\|(M+W)_{k,\le k-1}\|}{s_{(k-1)n'}(M_{\le k-1} + W_{\le k-1})}\right)^{-1} \\ \times \min\Big[s_{(k-1)n'}(M_{\le k-1} + W_{\le k-1}), s_{n'}\big(M_k + B_k\big) \Big] \end{aligned} \tag{5.78}$$

where we have abbreviated

$$B_k := W_k - (M+W)_{k,\leq k-1}(M_{\leq k-1}+W_{\leq k-1})^{-1}(M+W)_{\leq k-1,k} \tag{5.79}$$

Suppose that the event $\mathcal{E}_{k-1}(\beta_{k-1}(\gamma), W_{\leq k-1})$ holds. We condition on realizations of the submatrices $M_{\leq k-1}, M_{\leq k-1,k}, M_{k,\leq k-1}$ satisfying

$$\|M_{\leq k-1,k}\|, \|M_{k,\leq k-1}\| \leq (K_0+K)\sqrt{n} \tag{5.80}$$

and

$$s_{(k-1)n'}(M_{\leq k-1}+W_{\leq k-1}) \geq n^{-\beta_{k-1}(\gamma)}. \tag{5.81}$$

In particular,

$$\|(M+W)_{\leq k-1,k}\|, \|(M+W)_{k,\leq k-1}\| \leq 2n^{\gamma} \tag{5.82}$$

and by our restriction to $\mathcal{B}(K)$ and (5.36),

$$\|M_k + B_k\| \leq (K_0+K_1)\sqrt{n} + n^{\gamma} + 4n^{\gamma+\beta_{k-1}(\gamma)} \leq 6n^{\gamma+\beta_{k-1}(\gamma)}. \tag{5.83}$$

By Theorem 1.25 and independence of M_k from $M_{\leq k-1}, M_{k,\leq k-1}, M_{k,\leq k-1}$, there exists $\beta'_k = O(\gamma^2 + \beta_{k-1}(\gamma)^2)$ such that

$$\mathbb{P}\Big(s_{n'}(M_k+B_k) \leq n^{-\beta'_k}\Big) \ll_{K_0,K,\gamma,\delta,\sigma_0,\varepsilon,\mu_{4+\eta}} \sqrt{\frac{\log n}{n}}. \tag{5.84}$$

Restricting further to the event that this bound holds and substituting the above estimates into (5.78), we have

$$s_{kn'}(M_{\leq k}+W) \gg n^{-2\gamma-2\beta_{k-1}(\gamma)} \min(n^{-\beta_{k-1}(\gamma)}, n^{-\beta'_k}) \geq n^{-\beta_k(\gamma)} \tag{5.85}$$

for some $\beta_k(\gamma) = O(\gamma^2 + \beta_{k-1}(\gamma)^2)$. With this choice of $\beta_k(\gamma)$ we have shown

$$\mathbb{P}(\mathcal{E}_k(\beta_k(\gamma), W_{\leq k})^c \wedge \mathcal{E}_{k-1}(\beta_{k-1}(\gamma), W_{\leq k-1})) \ll_{K_0,K,\gamma,\delta,\sigma_0,\varepsilon,\mu_{4+\eta}} \sqrt{\frac{\log n}{n}}. \tag{5.86}$$

Applying this bound for all $2 \leq k' \leq k$ together with (5.76) and Bayes' rule we conclude that for any fixed k and any square matrix W of dimension at least kn' and operator norm at most n^{γ},

$$\mathbb{P}(\mathcal{E}_k(\beta_k(\gamma), W_{\leq k})^c) \ll_{K_0,K,\gamma,\delta,\sigma_0,\varepsilon,\mu_{4+\eta}} k\sqrt{\frac{\log n}{n}}. \tag{5.87}$$

The result now follows by taking $k = m$ and recalling that $m = O_{\varepsilon}(1)$.

Appendix A. Proofs of anti-concentration lemmas

In this appendix we proof Lemmas 2.5 and 2.7.

A.1. **Proof of Lemma 2.5.** (2.5) is immediate by our assumptions. It remains to show

$$\mathbb{E}\,|\mathrm{Re}(z\xi - w)|^2\,\mathbb{1}(|\xi| \le \kappa_0) \gg \frac{1}{\kappa_0}|\mathrm{Re}(z)|^2 \tag{A.1}$$

for all $z, w \in \mathbb{C}$ after rotating ξ by a phase if necessary. We may assume κ_0 is larger than any fixed constant. Let $\mathcal{E}$ denote the event $\{|\xi| \le \kappa_0\}$. By Chebyshev's inequality,

$$\mathbb{P}(\mathcal{E}) \ge 1 - \frac{1}{\kappa_0^2}. \tag{A.2}$$

Fix $z, w \in \mathbb{C}$. Write $\widetilde{\mathbb{E}} := \mathbb{E}(\cdot|\mathcal{E})$. By (A.2) and assuming κ_0 is sufficiently large we have that the left hand side of (A.1) is $\gg \widetilde{\mathbb{E}}|\mathrm{Re}(z\xi - w)|^2$, so it suffices to show

$$\widetilde{\mathbb{E}}|\mathrm{Re}(z\xi - w)|^2 \gg \frac{1}{\kappa_0}|\mathrm{Re}(z)|^2 \tag{A.3}$$

after rotating ξ by a phase. Denoting $\eta := \xi - \widetilde{\mathbb{E}}\xi$, we have

$$\widetilde{\mathbb{E}}|\mathrm{Re}(z\xi - w)|^2 = \widetilde{\mathbb{E}}|\mathrm{Re}(z\eta + (\widetilde{\mathbb{E}}\xi - w))|^2 = \widetilde{\mathbb{E}}|\mathrm{Re}(z\eta)|^2 + |\widetilde{\mathbb{E}}\xi - w|^2$$

so it suffices to show that after rotating ξ by a phase,

$$\widetilde{\mathbb{E}}|\mathrm{Re}(z\eta)|^2 \gg \frac{1}{\kappa_0}|\mathrm{Re}(z)|^2. \tag{A.4}$$

We first estimate the conditional variance of η. We have

$$\begin{aligned}
\widetilde{\mathbb{E}}|\eta|^2 &= \widetilde{\mathbb{E}}|\xi|^2 - |\widetilde{\mathbb{E}}\xi|^2 \\
&= \frac{1}{\mathbb{P}(\mathcal{E})}\,\mathbb{E}\,|\xi|^2\,\mathbb{1}_{\mathcal{E}} - \frac{1}{\mathbb{P}(\mathcal{E})^2}|\,\mathbb{E}\,\xi\,\mathbb{1}_{\mathcal{E}}\,|^2 \\
&= \frac{1}{\mathbb{P}(\mathcal{E})^2}\,\mathrm{Var}(\xi\,\mathbb{1}_{\mathcal{E}}) + \frac{1}{\mathbb{P}(\mathcal{E})}\left(1 - \frac{1}{\mathbb{P}(\mathcal{E})}\right)\mathbb{E}\,|\xi|^2\,\mathbb{1}_{\mathcal{E}} \\
&= \frac{1}{\mathbb{P}(\mathcal{E})^2}\left(\mathrm{Var}(\xi\,\mathbb{1}_{\mathcal{E}}) - \mathbb{P}(\mathcal{E}^c)\,\mathbb{E}\,|\xi|^2\,\mathbb{1}_{\mathcal{E}}\right) \\
&\gg \mathrm{Var}(\xi\,\mathbb{1}_{\mathcal{E}}) - O(1/\kappa_0^2)
\end{aligned}$$

where in the final line we applied (A.2), the assumption $\mathbb{E}\,|\xi|^2 = 1$, and assumed κ_0 is sufficiently large. Now by our assumption that ξ is κ_0-spread we have $\mathrm{Var}(\xi\,\mathbb{1}_{\mathcal{E}}) \gg 1/\kappa_0$, so

$$\widetilde{\mathbb{E}}|\eta|^2 \gg 1/\kappa_0 \tag{A.5}$$

taking κ_0 larger if necessary.

Now consider the covariance matrix

$$\Sigma_{\kappa_0} := \begin{pmatrix} \widetilde{\mathbb{E}}|\mathrm{Re}(\eta)|^2 & \widetilde{\mathbb{E}}(\mathrm{Re}(\eta)\mathrm{Im}(\eta)) \\ \widetilde{\mathbb{E}}(\mathrm{Re}(\eta)\mathrm{Im}(\eta)) & \widetilde{\mathbb{E}}|\mathrm{Im}(\eta)|^2 \end{pmatrix}. \tag{A.6}$$

Writing $z = a - ib$ and letting $x = (a \quad b)^{\mathsf{T}}$ be the associated column vector, we have

$$\widetilde{\mathbb{E}}|\mathrm{Re}(z\eta)|^2 = \widetilde{\mathbb{E}}|a\mathrm{Re}(\eta) + b\mathrm{Im}(\eta)|^2 = x^{\mathsf{T}}\Sigma_{\kappa_0}x. \tag{A.7}$$

Since Σ_{κ_0} has two non-negative eigenvalues $\sigma_1^2 \geq \sigma_2^2 \geq 0$ summing to $\widetilde{\mathbb{E}}|\eta|^2 \gg 1/\kappa_0$, it follows that $\sigma_1^2 \gg 1/\kappa_0$. We may rotate ξ by an appropriate phase to assume the corresponding eigenspace is spanned by $(1 \quad 0)^{\mathsf{T}}$. This gives

$$\widetilde{\mathbb{E}}|\mathrm{Re}(z\eta)|^2 \gg \sigma_1^2|\mathrm{Re}(z)|^2 \gg \frac{1}{\kappa_0}|\mathrm{Re}(z)|^2$$

as desired.

A.2. **Proof of Lemma 2.7.** We first need to recall a couple of lemmas from [TV10c, TV08].

Lemma A.1 (Fourier-analytic bound, cf. [TV10c, Lemma 6.1])**.** *Let ξ be a complex-valued random variable. For all $r > 0$ and any $v \in S^{n-1}$ we have*

$$p_{\xi,v}(r) \ll r^2 \int_{w\in\mathbb{C}:|w|\leq 1/r} \exp\Big(- c\sum_{j=1}^n \|wv_j\|_\xi^2\Big)\mathrm{d}w \tag{A.8}$$

where

$$\|z\|_\xi^2 := \mathbb{E}\,\|\mathrm{Re}(z(\xi - \xi'))\|_{\mathbb{R}/\mathbb{Z}}^2, \tag{A.9}$$

ξ' is an independent copy of ξ, and $\|x\|_{\mathbb{R}/\mathbb{Z}}$ denotes the distance from x to the nearest integer.

The next lemma gives an important property enjoyed by the "norm" $\|\cdot\|_\xi$ from Lemma A.1 under the assumption that ξ has κ-controlled second moment.

Lemma A.2 (cf. [TV08, Lemma 5.3])**.** *For any $\kappa > 0$ there are constants $c_1, c_2 > 0$ such that if ξ is κ-controlled, then $\|z\|_\xi \geq c_1|\mathrm{Re}(z)|$ whenever $|z| \leq c_2$.*

Now we begin the proof of Lemma 2.7. Let $r \geq 0$. We may assume $r \geq C_0\|v\|_\infty$ for any fixed constant $C_0 > 0$ depending only on κ. From Lemma A.1,

$$p_{\xi,v}(r) \ll r^2 \int_{|w|\leq 1/r} \exp\Big(- c\sum_{j=1}^n \|wv_j\|_\xi^2\Big)\mathrm{d}w.$$

If C_0 is sufficiently large depending on κ, it follows from Lemma A.2 that whenever $|w| \leq 1/r$, $\|wv_j\|_\xi \geq c_1|\mathrm{Re}(wv_j)|$, giving

$$p_{\xi,v}(r) \ll r^2 \int_{|w|\leq 1/r} \exp\Big(- c'\sum_{j=1}^n (\mathrm{Re}(wv_j))^2\Big)\mathrm{d}w$$

where c' depends only on κ. By change of variable,

$$p_{\xi,v}(r) \ll \int_{|w|\leq 1} \exp\Big(- \frac{c'}{r^2}\sum_{j=1}^n (\mathrm{Re}(wv_j))^2\Big)\mathrm{d}w. \tag{A.10}$$

Write $v_j = r_j e^{i\theta_j}$ for each $j \in [n]$. Since $v \in S^{n-1}$ we have $\sum_{j=1}^n r_j^2 = 1$. By Jensen's inequality,

$$p_{\xi,v}(r) \ll \int_{|w|\le 1} \exp\Big(-\frac{c'}{r^2} \sum_{j=1}^n r_j^2 \big(\mathrm{Re}(we^{i\theta_j})\big)^2 \Big) \mathrm{d}w$$
$$\le \int_{|w|\le 1} \sum_{j=1}^n r_j^2 \exp\Big(-\frac{c'}{r^2} \big(\mathrm{Re}(we^{i\theta_j})\big)^2 \Big) \mathrm{d}w.$$

By rotational invariance the last expression is equal to

$$\sum_{j=1}^n r_j^2 \int_{|w|\le 1} \exp\Big(-\frac{c'}{r^2} (\mathrm{Re}(w))^2 \Big) \mathrm{d}w = \int_{|w|\le 1} \exp\Big(-\frac{c'}{r^2} (\mathrm{Re}(w))^2 \Big) \mathrm{d}w$$

which by direct computation is seen to be of size $O(r)$ (with implied constant depending on κ). Together with our assumption that $r \ge C_0 \|v\|_\infty$ this gives (2.7).

References

[AGZ10] Greg W. Anderson, Alice Guionnet, and Ofer Zeitouni. *An introduction to random matrices*, volume 118 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2010.

[Ara80] T. V. Arak. Approximation of n-fold convolutions of distributions, having a nonnegative characteristic function, with accompanying laws. *Teor. Veroyatnost. i Primenen.*, 25(2):225–246, 1980.

[Ara81] T. V. Arak. On the rate of convergence in Kolmogorov's uniform limit theorem. *Teor. Veroyatnost. i Primenen.*, 26(2):225–245, 1981.

[ARS15] Johnatan Aljadeff, David Renfrew, and Merav Stern. Eigenvalues of block structured asymmetric random matrices. *J. Math. Phys.*, 56(10):103502, 14, 2015.

[AS04] Noga Alon and Asaf Shapira. Testing subgraphs in directed graphs. *J. Comput. System Sci.*, 69(3):353–382, 2004.

[BC12] Charles Bordenave and Djalil Chafaï. Around the circular law. *Probab. Surv.*, 9:1–89, 2012.

[BEYY] Paul Bourgade, Laszlo Erdos, Horng-Tzer Yau, and Jun Yin. Universality for a class of random band matrices. Preprint available at arXiv:1602.02312.

[BSY88] Z. D. Bai, Jack W. Silverstein, and Y. Q. Yin. A note on the largest eigenvalue of a large-dimensional sample covariance matrix. *J. Multivariate Anal.*, 26(2):166–168, 1988.

[BT87] J. Bourgain and L. Tzafriri. Invertibility of "large" submatrices with applications to the geometry of Banach spaces and harmonic analysis. *Israel J. Math.*, 57(2):137–224, 1987.

[BvH16] Afonso S. Bandeira and Ramon van Handel. Sharp nonasymptotic bounds on the norm of random matrices with independent entries. *Ann. Probab.*, 44(4):2479–2506, 2016.

[CHNR] Nicholas A. Cook, Walid Hachem, Jamal Najim, and David Renfrew. Limiting spectral distribution for non-hermitian random matrices with a variance profile. In preparation.

[Coo16] Nicholas A. Cook. *Spectral properties of non-Hermitian random matrices*. PhD thesis, University of California, Los Angeles, 2016.

[Ede88] Alan Edelman. Eigenvalues and condition numbers of random matrices. *SIAM J. Matrix Anal. Appl.*, 9(4):543–560, 1988.

[Ede89] Alan Stuart Edelman. *Eigenvalues and condition numbers of random matrices*. PhD thesis, Massachusetts Institute of Technology, 1989.

[GEZ16] F. Gëttse, Yu. S. Eliseeva, and A. Yu. Zaĭtsev. Arak inequalities for concentration functions and the Littlewood-Offord problem. *Dokl. Akad. Nauk*, 467(5):514–518, 2016.

[Gow97] W. T. Gowers. Lower bounds of tower type for Szemerédi's uniformity lemma. *Geom. Funct. Anal.*, 7(2):322–337, 1997.

[HLN07] Walid Hachem, Philippe Loubaton, and Jamal Najim. Deterministic equivalents for certain functionals of large random matrices. *Ann. Appl. Probab.*, 17(3):875–930, 2007.

[Kom] János Komlós. Circulated manuscript, 1977. Edited version available online at: http://www.math.rutgers.edu/∼komlos/01short.pdf.

[Kom67] J. Komlós. On the determinant of (0, 1) matrices. *Studia Sci. Math. Hungar*, 2:7–21, 1967.

[Kom68] J. Komlós. On the determinant of random matrices. *Studia Sci. Math. Hungar.*, 3:387–399, 1968.

[KS96] J. Komlós and M. Simonovits. Szemerédi's regularity lemma and its applications in graph theory. In *Combinatorics, Paul Erdős is eighty, Vol. 2 (Keszthely, 1993)*, volume 2 of *Bolyai Soc. Math. Stud.*, pages 295–352. János Bolyai Math. Soc., Budapest, 1996.

[Lat05] Rafał Latała. Some estimates of norms of random matrices. *Proc. Amer. Math. Soc.*, 133(5):1273–1282 (electronic), 2005.

[LLT+17] Alexander E. Litvak, Anna Lytova, Konstantin Tikhomirov, Nicole Tomczak-Jaegermann, and Pierre Youssef. Adjacency matrices of random digraphs: singularity and anti-concentration. *J. Math. Anal. Appl.*, 445(2):1447–1491, 2017.

[LP86] Françoise Lust-Piquard. Inégalités de Khintchine dans C_p $(1 < p < \infty)$. *C. R. Acad. Sci. Paris Sér. I Math.*, 303(7):289–292, 1986.

[LPP91] Françoise Lust-Piquard and Gilles Pisier. Noncommutative Khintchine and Paley inequalities. *Ark. Mat.*, 29(2):241–260, 1991.

[LPR+05] A. E. Litvak, A. Pajor, M. Rudelson, N. Tomczak-Jaegermann, and R. Vershynin. Euclidean embeddings in spaces of finite volume ratio via random matrices. *J. Reine Angew. Math.*, 589:1–19, 2005.

[LPRTJ05] A. E. Litvak, A. Pajor, M. Rudelson, and N. Tomczak-Jaegermann. Smallest singular value of random matrices and geometry of random polytopes. *Adv. Math.*, 195(2):491–523, 2005.

[LR12] Alexander E. Litvak and Omar Rivasplata. Smallest singular value of sparse random matrices. *Studia Math.*, 212(3):195–218, 2012.

[MSS14] A.W. Marcus, D.A. Spielman, and N. Srivastava. Ramanujan graphs and the solution of the kadison–singer problem. In *Proc. ICM, Vol III*, pages 375–386, 2014.

[NV13] Hoi H. Nguyen and Van H. Vu. Small ball probability, inverse theorems, and applications. In *Erdős centennial*, volume 25 of *Bolyai Soc. Math. Stud.*, pages 409–463. János Bolyai Math. Soc., Budapest, 2013.

[NV16] Hoi H. Nguyen and Van H. Vu. Normal vector of a random hyperplane. Preprint available at arXiv:1604.04897, 04 2016.

[Oli10] Roberto Imbuzeiro Oliveira. Sums of random Hermitian matrices and an inequality by Rudelson. *Electron. Commun. Probab.*, 15:203–212, 2010.

[RA06] Kanaka Rajan and LF Abbott. Eigenvalue spectra of random matrices for neural networks. *Physical review letters*, 97(18):188104, 2006.

[RT] Elizaveta Rebrova and Konstantin Tikhomirov. Covering of random ellipsoids, and invertibility of matrices with i.i.d. heavy-tailed entries. Preprint available at arXiv:1508.06690.

[Rud08] Mark Rudelson. Invertibility of random matrices: norm of the inverse. *Ann. of Math. (2)*, 168(2):575–600, 2008.

[RV08a] Mark Rudelson and Roman Vershynin. The least singular value of a random square matrix is $O(n^{-1/2})$. *C. R. Math. Acad. Sci. Paris*, 346(15-16):893–896, 2008.

[RV08b] Mark Rudelson and Roman Vershynin. The Littlewood-Offord problem and invertibility of random matrices. *Adv. Math.*, 218(2):600–633, 2008.

[RZ16] Mark Rudelson and Ofer Zeitouni. Singular values of Gaussian matrices and permanent estimators. *Random Structures Algorithms*, 48(1):183–212, 2016.

[Seg00] Yoav Seginer. The expected norm of random matrices. *Combin. Probab. Comput.*, 9(2):149–166, 2000.

[SS12] Daniel A. Spielman and Nikhil Srivastava. An elementary proof of the restricted invertibility theorem. *Israel J. Math.*, 190:83–91, 2012.

[SST06] Arvind Sankar, Daniel A. Spielman, and Shang-Hua Teng. Smoothed analysis of the condition numbers and growth factors of matrices. *SIAM J. Matrix Anal. Appl.*, 28(2):446–476 (electronic), 2006.

[Sze78] Endre Szemerédi. Regular partitions of graphs. In *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*, volume 260 of *Colloq. Internat. CNRS*, pages 399–401. CNRS, Paris, 1978.

[Tal96] Michel Talagrand. A new look at independence. *Ann. Probab.*, 24(1):1–34, 1996.

[Tro12] Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.*, 12(4):389–434, 2012.

[TV04] Antonia M Tulino and Sergio Verdú. *Random matrix theory and wireless communications*, volume 1. Now Publishers Inc, 2004.

[TV07] Terence Tao and Van H. Vu. On the singularity probability of random Bernoulli matrices. *J. Amer. Math. Soc.*, 20(3):603–628, 2007.

[TV08] Terence Tao and Van H. Vu. Random matrices: the circular law. *Commun. Contemp. Math.*, 10(2):261–307, 2008.

[TV09] Terence Tao and Van H. Vu. Inverse Littlewood-Offord theorems and the condition number of random discrete matrices. *Ann. of Math. (2)*, 169(2):595–632, 2009.

[TV10a] Terence Tao and Van Vu. Random matrices: the distribution of the smallest singular values. *Geom. Funct. Anal.*, 20(1):260–297, 2010.

[TV10b] Terence Tao and Van H. Vu. Random matrices: universality of ESDs and the circular law. *Ann. Probab.*, 38(5):2023–2065, 2010. With an appendix by Manjunath Krishnapur.

[TV10c] Terence Tao and Van H. Vu. Smooth analysis of the condition number and the least singular value. *Math. Comp.*, 79(272):2333–2352, 2010.

[Ver11] Roman Vershynin. Spectral norm of products of random and deterministic matrices. *Probab. Theory Related Fields*, 150(3-4):471–509, 2011.

[vH] Ramon van Handel. On the spectral norm of gaussian random matrices. Preprint available at arXiv:1502.05003.

[vNG47] John von Neumann and H. H. Goldstine. Numerical inverting of matrices of high order. *Bull. Amer. Math. Soc.*, 53:1021–1099, 1947.

[YBK88] Y. Q. Yin, Z. D. Bai, and P. R. Krishnaiah. On the limit of the largest eigenvalue of the large-dimensional sample covariance matrix. *Probab. Theory Related Fields*, 78(4):509–521, 1988.

Department of Mathematics, Stanford University
E-mail address: nickcook@stanford.edu